arXiv:1512.01184v1 [math.GT] 3 Dec 2015

GRAPH COBORDISMS AND HEEGAARD FLOER HOMOLOGY

IAN ZEMKE

Abstract. In this paper we construct graph cobordism maps for all flavors of Heegaard Floer homology and prove invariance of such maps. In the process we also finish the proof of invariance of the 4-manifold invariants constructed by Ozsváth and Szabó. The graph cobordism maps are defined for a very general class of cobordisms, including cobordisms which are disconnected, or have disconnected or empty ends. We also define a sort of TQFT for graphs inside a fixed 3-manifold. As an application we compute the chain homotopy type of the π_1-action on CF°, showing that it is chain homotopic to the identity map for the $+,-$ and ∞ flavors, but the chain homotopy is not U-equivariant. For the cobordisms originally considered by Ozsváth and Szabó (connected path cobordisms with connected ends and a single path), this implies that, after taking homology, the cobordism maps don't depend on the path for the $+,-$ and ∞ flavors.

Contents

1. Introduction 2
2. Acknowledgements 5
3. Organization of the paper 5
4. Preliminaries on Heegaard Floer homology 6
4.1. Holomorphic disks in the cylindrical setup 6
4.2. The colored Heegaard Floer complexes 7
4.3. Holomorphic triangle maps 8
4.4. Acceptable domains and combinatorial Maslov indices 9
4.5. Admissibility 11
4.6. Heegaard moves and naturality of the 3-manifold invariants 14
5. Homological actions on the Heegaard Floer groups 15
5.1. Construction of the relative homology actions 15
5.2. Naturality of the relative homology maps 17
6. Free stabilization maps 23
6.1. Gluing data, neck stretching, and gluing moduli spaces 23
6.2. Counting holomorphic strips after free stabilization 25
6.3. Dependence on gluing data 27
6.4. Free stabilization of the triangle maps 31
6.5. Topological preliminaries about based embeddings of Heegaard surfaces 35
6.6. Invariance of free stabilization maps 35
7. Graph action map 37
8. 1- and 3-handles 46
8.1. Preliminaries on 1- and 3-handles 46
8.2. Gluing data for 1- and 3-handle maps 47
8.3. Counting holomorphic disks after surgery on $\mathbb{S}_0$ 48
8.4. Counting holomorphic triangles after surgery on $\mathbb{S}_0$ 50
8.5. Dependence on gluing data 51
8.6. Invariance of 1- and 3-handle maps 52
9. 2-handle maps 52
9.1. Maps associated to a framed link 52
9.2. (1,2)-stabilization and triangle maps 54
9.3. Invariance of maps associated to a framed link 57
10. Invariance of path cobordism maps 58

11. Construction and invariance of the graph cobordism maps when $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$ 62
11.1. 0- and 4-handle maps 62
11.2. Main construction 62
11.3. Invariance of the graph cobordism maps from $\widetilde{\Gamma}_0$ 65
11.4. Invariance of the cobordism maps from $N(\Gamma_0)$ 66
11.5. Invariance of the graph cobordism maps from subdivision 67
11.6. Dependence on cyclic ordering 67
12. The composition law for $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$ 68
13. Cobordisms with $\pi_0(Y_i) \not\twoheadrightarrow \pi_0(W)$ 75
13.1. Main technical lemma 75
13.2. Construction of graph cobordism maps for arbitrary cobordisms 76
13.3. Composition law for arbitrary cobordisms 77
14. Moving basepoints, the π_1 action, and the relation between the graph and path cobordism maps 79
14.1. Transition map computation 79
14.2. An easier model computation 87
14.3. Basepoint moving map, the $\pi_1(Y,p)$-action, and the relation between path and graph cobordism maps 88
References 91

1. Introduction

Heegaard Floer homology, introduced by Ozsváth and Szabó associates functorial invariants to 3- and 4-manifolds. To an oriented, closed 3-manifold with a collection of basepoints, $(Y, \mathbf{w})$, and a Spin^c structure $\mathfrak{s}$, one associates homology groups

$$HF^\circ(Y, \mathbf{w}, \mathfrak{s})$$

for $\circ = +, -, \infty, \widehat{\ }$. These are modules over the ring $\mathbb{Z}_2[U_{w_1}, \dots, U_{w_n}]$ where $\mathbf{w} = \{w_1, \dots, w_n\}$. Originally defined only for connected 3-manifolds Y^3, in this paper we consider arbitrary closed and oriented Y which have at least one basepoint in each component. Ozsváth and Szabó showed that the 3-manifold invariants fit into the framework of a (3+1)-dimensional TQFT:

Theorem 1.1 ([OS06])**.** *Given a connected 4-manifold W with connected ends, (Y_1, w_1) and (Y_2, w_2), and a 4-dimensional* Spin^c *structure $\mathfrak{s}$ on W, there are maps*

$$F^\circ_{W,\mathfrak{s},A} : HF^\circ(Y_1, w_1, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, w_2, \mathfrak{s}|_{Y_2})$$

for some choice extra data A.

The dependence on the extra data A was originally not understood, but it has been widely believed that the minimal amount of extra data needed is a path in W from w_1 to w_2. There have been many steps towards a proof of invariance for such maps, notably [OS06], [JT12], and [Juh15]. In this paper we provide the final steps towards proving invariance:

Theorem A. *Given a cobordism $(W, \boldsymbol{\gamma}) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where $\boldsymbol{\gamma}$ is a collection of paths from the basepoints in $\mathbf{w}_1$ to the basepoints in $\mathbf{w}_2$, there is a map*

$$F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} : HF^\circ(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathbf{w}_2, \mathfrak{s}|_{Y_2}),$$

which is an invariant. Here W, Y_1 and Y_2 are allowed to be disconnected. The maps are $\mathbb{Z}_2[U_{\mathbf{w}}]$-module homomorphisms where we identify two variables U_{w_1} and U_{w_2} if there is a path in $\boldsymbol{\gamma}$ from $w_1 \in \mathbf{w}_1$ to $w_2 \in \mathbf{w}_2$.

In [OS06], for a closed, connected 4-manifold X with Spin^c structure $\mathfrak{s}$, Ozsváth and Szabó construct a map

$$\Phi_{X,\mathfrak{s}} : \mathbb{Z}_2[U] \otimes \Lambda^*(H_1(X)/\operatorname{Tors}) \to \mathbb{Z}_2.$$

To define such a map, one must choose two balls B_1, B_2 in X, an admissible cut $N \subseteq X$, as well as paths from a point $n \in N$ to ∂B_1 and ∂B_2. Given an admissible cut $N \subseteq X$, any two configurations of the above

form are diffeomorphic, and hence Theorem A immediately shows that the untwisted version of this map, i.e. the map $\Phi_{X,\mathfrak{s}} : \mathbb{Z}_2[U] \to \mathbb{Z}_2$, depends only on N. In [OS06], it is shown to be independent of the admissible cut N.

In [Zem15] the author defines a "graph TQFT" for the hat Heegaard Floer homology using the sutured Floer homology TQFT developed by Andás Juhász in [Juh06a] and [Juh15]. To a 4-manifold cobordism with embedded graph $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ such that $\Gamma \cap Y_i = \mathbf{w}_i$, the $\widehat{HF}$ graph TQFT associates a map

$$\widehat{F}_{W,\Gamma} : \widehat{HF}(Y_1, \mathbf{w}_1) \to \widehat{HF}(Y_2, \mathbf{w}_2),$$

which decomposes over Spin^c structures and is functorial under composition of graph cobordisms. The $\widehat{HF}$ graph TQFT defined in this manner allows for more cobordisms than in the original construction by Ozsváth and Szabó. It allows for disconnected 4-manifolds W and for disconnected or empty 3-manifolds Y_i as ends of 4-manifolds. The extra flexibility comes at the expense of keeping track of an embedded graph $\Gamma \subseteq W$. It is known that the $\pi_1(Y, p)$ action does not vanish on $\widehat{HF}(Y, p)$, and hence such a TQFT cannot be constructed without the additional data of a graph or collection of paths. There is a version of sutured Floer homology that generalizes HF^- for closed manifolds, developed by Alishahi and Eftekhary in [AE11], though we will use other techniques to construct a graph TQFT for the $+, -$ and ∞ flavors of Heegaard Floer homology.

In [Zem15], we showed that the graph cobordism maps in cylindrical cobordisms $W = Y \times I$ determine an action of $\Lambda^*(H_1(Y, \mathbf{w}; \mathbb{Z})/\operatorname{Tors})$ on $\widehat{HF}(Y, \mathbf{w})$, as well as maps corresponding to adding or removing a basepoint. Those maps, together with the path cobordism maps, uniquely determine the graph cobordism maps. In this paper we construct maps corresponding to adding or removing basepoints (which we call the free stabilization maps), and maps corresponding to a new relative homology action, and use these together with maps corresponding to handle attachment to define graph cobordism maps. To a path λ between two basepoints w_1 and w_2 which is embedded in the underlying Heegaard surface of a diagram $\mathcal{H}$ of $(Y, \mathbf{w})$, we define a map

$$A_\lambda : CF^\circ(\mathcal{H}, \mathfrak{s}) \to CF^\circ(\mathcal{H}, \mathfrak{s})$$

which satisfies

$$A_\lambda \partial + \partial A_\lambda = U_{w_1} + U_{w_2}.$$

The situation is similar, but slightly different, for other Floer theories. For example in embedded contact homology (cf. [HT09, Sec. 2.5]), a complex $C_*(Y, \lambda, \Gamma)$ is defined independently of a basepoint and the U map is defined by first picking a basepoint w and counting certain holomorphic curves which depend on the basepoint w to get a map U_w. A path from w to w' yields a chain homotopy between the U_w and $U_{w'}$ maps on the complex $C_*(Y, \lambda, \Gamma)$. The situation is similar for the U-map in the monopole Floer homology theory developed by Kronheimer and Mrowka (cf. [KM07]).

Several novel features appear for our new graph TQFT for HF° which are not present in the hat flavor of the graph TQFT constructed in [Zem15]. Firstly we must now decorate the underlying graph in a graph cobordism with a ribbon structure, i.e. a cyclic ordering of the edges adjacent to each vertex. In the graph TQFT for $\widehat{HF}$ constructed previously, the maps depended only on the isotopy class of a regular neighborhood of the graph. In addition, Heegaard Floer homology has a module structure over a polynomial ring, which acts trivially over the hat flavor. To define functorial maps, we must work with colored sets of basepoints and colored graphs. If $\mathbf{w}$ is a set of basepoints, a coloring $(\sigma, \mathfrak{P})$ of $\mathbf{w}$ is a set $\mathfrak{P}$ indexing a collection of formal variables, as well as a map $\sigma : \mathbf{w} \to \mathfrak{P}$. If Γ is a graph, a coloring $(\sigma, \mathfrak{P})$ of Γ consists of a set $\mathfrak{P}$ and a map $\sigma : C(\Gamma) \to \mathfrak{P}$, where $C(\Gamma)$ denotes the set of connected components of Γ. We let $\mathbb{Z}_2[U_\mathfrak{P}]$ denote the free abelian alegebra over $\mathbb{Z}_2$ and the variables U_p for $p \in \mathfrak{P}$. The colored Heegaard Floer homology groups

$$HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s})$$

are then defined by identifying the variable U_w with $U_{\sigma(w)}$ before taking homology. They are $\mathbb{Z}_2[U_\mathfrak{P}]$ modules. Variables corresponding to colors in $\mathfrak{P}$ which are not represented by any basepoints act essentially by extension of scalars. We will usually omit writing $\mathfrak{P}$ and write $HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{s})$ (or just $HF^\circ(Y, \mathbf{w}, \mathfrak{s})$) unless we need to emphasize the set $\mathfrak{P}$ or the coloring σ.

As a first step towards the full graph cobordism TQFT, we define a sort TQFT for graphs inside of a fixed, closed 3-manifold Y^3. We say that a graph Γ is a "flow graph" from vertices V_0 to V_1 if each vertex in V_i has valence one in Γ and $V_0 \cap V_1 = \varnothing$. We use the functorial notation $\Gamma : V_0 \to V_1$. Notice that if

$\Gamma : V_0 \to V_1$ is a flow graph, and $(\sigma, \mathfrak{P})$ is a coloring of Γ, then $(\sigma, \mathfrak{P})$ restricts to a coloring of V_0 and V_1. An example of a flow graph in a 3-manifold Y is shown in Figure 1. We prove the following:

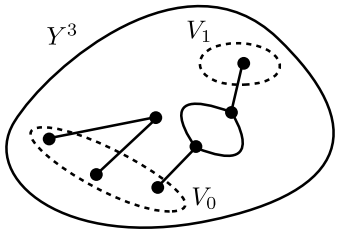

FIGURE 1. A flow graph $\Gamma : V_0 \to V_1$, embedded in a 3-manifold Y^3. The dashed ovals indicate the vertex sets V_0 and V_1.

Theorem B. *Given a ribbon flow graph $\Gamma : V_0 \to V_1$ embedded in Y^3 with coloring $(\sigma, \mathfrak{P})$, and* Spin^c *structure $\mathfrak{s}$ on Y, there is a morphism of $\mathbb{Z}_2[U_{\mathfrak{P}}]$-modules*

$$A_{\Gamma,\mathfrak{s}} : HF^\circ(Y, V_0, \sigma|_{V_0}, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, V_1, \sigma|_{V_1}, \mathfrak{P}, \mathfrak{s})$$

which is functorial under concatenating graph components together.

We call the map $A_{\Gamma,\mathfrak{s}}$ the **graph action map**, since it encodes the $\Lambda^*(H_1(Y,\mathbb{Z})/\operatorname{Tors})$ action on HF°. In the case that $\Gamma : V_0 \to V_1$ consists of a collection of arcs from a point in V_0 to a point in V_1, the map $A_{\Gamma,\mathfrak{s}}$ is the map induced by the diffeomorphism corresponding to finger moves along the arcs in Γ.

We use the maps $A_{\Gamma,\mathfrak{s}}$ to define a full graph TQFT for cobordisms with embedded, colored ribbon graphs:

Theorem C. *Given a graph cobordism $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where Γ is a ribbon graph with coloring $(\sigma, \mathfrak{P})$, there is a homomorphism of $\mathbb{Z}_2[U_{\mathfrak{P}}]$-modules*

$$F^\circ_{W,\Gamma,\mathfrak{s}} : HF^\circ(Y_1, \mathbf{w}_1, \sigma|_{\mathbf{w}_1}, \mathfrak{P}, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathbf{w}_2, \sigma|_{\mathbf{w}_2}, \mathfrak{P}, \mathfrak{s}|_{Y_2}).$$

These are invariants of the cobordism with embedded colored ribbon graph.

In the above theorem, the cobordism W can be disconnected, and the ends Y_i can be disconnected or even emtpy. There are also no restrictions on the embedded ribbon graph Γ other than that the vertices of the graph lying in Y_i must have valence 1. The new graph cobordism maps coincide with the path cobordism maps when the graph is a path:

Theorem D. *If $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a graph cobordism and Γ is a collection of paths $\boldsymbol{\gamma}$ from the basepoints in Y_1 to the basepoints in Y_2, then*

$$F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}},$$

where the map on the left is the graph cobordism map, and the map on the right is the path cobordism map.

The graph cobordism maps essentially satisfy the standard composition law. The only caveat is that one must be caseful about the coloring. In the following we will write $F^\circ_{W,\Gamma,\sigma,\mathfrak{P},\mathfrak{s}}$ to emphasize the coloring:

Theorem E. *Suppose that (W, Γ) is ribbon graph cobordism with coloring $(\sigma, \mathfrak{P})$. Suppose that (W, Γ) decomposes as $(W, \Gamma) = (W_2, \Gamma_2) \cup (W_1, \Gamma_1)$. If $\mathfrak{s}_i$ is a* Spin^c *structure on W_i, then we have the following:*

$$F^\circ_{W_2,\Gamma_2,\sigma|_{\Gamma_2},\mathfrak{P},\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma_1,\sigma|_{\Gamma_1},\mathfrak{P},\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_2}=\mathfrak{s}_2 \\ \mathfrak{s}|_{W_1}=\mathfrak{s}_1}} F^\circ_{W,\Gamma,\sigma,\mathfrak{P},\mathfrak{s}}.$$

Notice that all of the maps in the above theorem are maps of $\mathbb{Z}_2[U_{\mathfrak{P}}]$-modules.

It is interesting to compare the cobordism maps in Heegaard Floer homology to the cobordism maps in monopole Floer homology constructed in [KM07]. Due to the work of Kutluhan, Lee and Taubes in [KLT10a],[KLT10b],[KLT10c],[KLT11], and [KLT12], and due to Colin, Ghiggini and Honda in [CGH10],

[CGH12a],[CGH12b] and [CGH12c], we now know that certain versions of Heegaard Floer homology, monopole Floer homology, and embedded contact homology are isomorphic. To the authors knowledge, a proof that these isomorphisms extend to the level of 4-manifold cobordism maps in Heegaard Floer homology and monopole Floer homology has not appeared in the literature, though the two theories are strongly suspected to be equivalent. In light of this, it is natural to ask how the path cobordism maps in Heegaard Floer homology depend on the choices of paths, since one does not need to decorate cobordisms with paths in monopole Floer homology.

To this end, we show in this paper that the 3-dimensional graph action maps yield convenient expressions for the effect of moving basepoints along paths, and we are able to perform enough model computations to compute explicitly the $\mathbb{Z}_2[U_w]$-chain homotopy type of the $\pi_1(Y,w)$-action on all flavors of Heegaard Floer homology.

Theorem F. *If $(Y,\mathbf{w})$ is a 3-manifold and $w\in\mathbf{w}$, then the $\pi_1(Y,w)$-action has the $\mathbb{Z}_2[U_{\mathbf{w}}]$-chain homotopy type*

$$\gamma_* = 1+\Phi_w[\gamma]$$

on $CF^\circ(Y,\mathbf{w},\mathfrak{s})$ for $\circ\in\{+,-,\infty\}$, where Φ_w is the map

$$\Phi_w=\sum_{i=1}^{\infty} n\partial^n U_w^{n-1},$$

where $\partial=\sum_{i=0}^{\infty}\partial^n U_w^n$ and $[\gamma]$ denotes the $H_1(Y,\mathbb{Z})/\operatorname{Tors}$ action. The map Φ_w is chain homotopic to 0, though the chain homotopy isn't $U_{\mathbf{w}}$-equivariant. Hence γ_* is chain homotopic to 1, though not $U_{\mathbf{w}}$-equivariantly so.

As a consequence, we have the following result:

Theorem G. *If $(W,\boldsymbol{\gamma})$ is path cobordism where all the paths have distinct colors (for example if there is only one path), then the map*

$$F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}}: HF^\circ(Y_1,\mathbf{w}_1,\mathfrak{s}|_{Y_1})\to HF^\circ(Y_2,\mathbf{w}_2,\mathfrak{s}|_{Y_2})$$

is independent of the choice of paths $\boldsymbol{\gamma}$ for $\circ=+,-,\infty$.

The stronger invariant, the map on chain complexes up to equivariant chain homotopy, does depend on the choice of paths–but after taking homology the dependence vanishes.

2. Acknowledgements

I would like to thank my advisor, Ciprian Manolescu, for insightful discussions and useful advice, as well as Matthew Stoffregen, Mike Miller and Yajing Liu for helpful conversations. I would like to thank Ciprian Manolescu and Peter Ozsváth for sharing with me an outline of how the triangle maps interacted with the 1- and 3-handle maps. I would also like to thank Robert Lipshitz for explaining some of the subtleties of stabilization invariance to me.

3. Organization of the paper

The organization of the paper is as follows:

- After discussing background material on Heegaard Floer homology in Section 4, we first define relative homology maps on the Heegaard Floer homology groups in Section 5. These are maps associated to closed loops or paths between two basepoints in Y, generalizing the action of $\Lambda^*(H_1(Y;\mathbb{Z})/\operatorname{Tors})$ on $HF^\circ(Y,w,\mathfrak{s})$ in [OS04b]. To a closed loop γ, we get a chain map A_γ, and to a path λ from w_1 to w_2, we get a chain homotopy A_λ between U_{w_1} and U_{w_2}.
- In Section 6 we define maps corresponding to adding or removing a basepoint, which we call free stabilization maps. We denote the free stabilization maps by S^+_w and S^-_w, depending on whether we are adding or removing the basepoint w.
- In Section 7 we then construct the graph action map associated to a ribbon flow graph embedded in a 3-manifold using the relative homology maps, and the free stabilization maps. In this section we prove Theorem B.

- In Sections 8-10, we define 1-,2-, and 3-handle maps, show that they are invariants, and show that they fit into the framework developed in [Juh15] and hence determine cobordism maps for cobordisms with embedded paths between the basepoints, proving Theorem A. Our contribution to this invariance result is mainly the gluing lemmas necessary to show that the 1-,2-, and 3-handle maps are invariants.
- In Section 11 we construct the graph cobordism maps for graph cobordisms $(W,\Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where $\pi_0(Y_i) \hookrightarrow \pi_0(W)$ is a surjection for $i = 1, 2$. We show that they are invariant, proving a weak version of Theorem C. In Section 12 we show that the composition law holds for the cobordisms with $\pi_0(Y_i) \hookrightarrow \pi_0(W)$ a surjection, proving a weaker version of Theorem E. Finally in Section 13 we construct maps for cobordisms where $\pi_0(Y_i) \hookrightarrow \pi_0(W)$ is not a surjection. We then show that they are invariants, and that they satisfy the composition law, proving Theorems C and E in full generality.
- Finally in Section 14 we consider the effect of moving basepoints. Through a challenging model computation, we show that the effect of moving basepoints is captured by the graph action map. As a consequence, we arrive at a proof of Theorem D, showing that the path cobordism maps agree with the graph cobordism maps when the graph is a collection of paths. We also are able to explicitly compute the π_1 action on CF° and thus prove Theorem F. By exhibiting a chain homotopy which is not $U_{\mathbf{w}}$-equivariant, we arrive at Theorem G.

4. Preliminaries on Heegaard Floer homology

We begin with the construction and definition of the 3-manifold invariants. A multibased 3-manifold $(Y, \mathbf{w})$ is a 3-manifold Y with a finite collection of basepoints $\mathbf{w} \subseteq Y$ such that $\pi_0(\mathbf{w}) \hookrightarrow \pi_0(Y)$ is a surjection.

Definition 4.1. *A **multipointed Heegaard diagram** $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ for $(Y, \mathbf{w})$ is a collection satisfying the following:*

(1) $\Sigma \subseteq Y$ is a closed, oriented surface of genus g embedded in Y such that $\mathbf{w} \subseteq \Sigma$;
(2) $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_{g+\ell-1})$ is a collection of disjoint, simple closed curves on Σ bounding compressing disks in Y, where $|\mathbf{w}| = \ell$;
(3) $\boldsymbol{\beta} = (\beta_1, \dots, \beta_{g+\ell-1})$ is another collection of disjoint, simple closed curves in Σ bounding compressing disks in Y;
(4) the collections $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ each span a g-dimensional sublattice of $H_1(\Sigma, \mathbb{Z})$;
(5) each α_i and β_j intersect transversely.

We will also require that our diagrams satisfy admissibility requirements, which we will discuss in Section 4.5. Before we define the complexes precisely, we need to discuss the cylindrical formulation of Heegaard Floer homology.

4.1. Holomorphic disks in the cylindrical setup.

In the original formulation of Heegaard Floer homology in [OS04b] and [OS06], one defines the flavors of Heegaard Floer homology as modified versions of Lagrangian intersection Floer homology of the tori

$$\mathbb{T}_\alpha = \alpha_1 \times \cdots \times \alpha_{g+\ell-1}, \quad \text{and} \quad \mathbb{T}_\beta = \beta_1 \times \cdots \times \beta_{g+\ell-1}$$

inside of $\mathrm{Sym}^{g+\ell-1}(\Sigma)$. An equivalent construction, the cylindrical formulation, was given by Robert Lipshitz in [Lip06]. Throughout the paper we will use this formulation, mainly because gluing lemmas are considerably easier to prove.

Following [Lip06], we define $W_{\mathbb{D}} = \Sigma \times I \times \mathbb{R}$ (where $I = [0, 1]$) and define $\pi_{\mathbb{D}} : W_{\mathbb{D}} \to I \times \mathbb{R}$, $\pi_{\mathbb{R}} : W_{\mathbb{D}} \to \mathbb{R}$ and $\pi_\Sigma : W_{\mathbb{D}} \to \Sigma$ to be the projection maps. We define $C_\alpha = \boldsymbol{\alpha} \times \{1\} \times \mathbb{R}$ and $C_\beta = \boldsymbol{\beta} \times \{0\} \times \mathbb{R}$. We consider certain almost complex structures J which are perturbations of a split almost complex structure $\mathfrak{j}_\Sigma \times \mathfrak{j}_{\mathbb{D}}$ (see [Lip06] for the precise requirements).

We consider Riemann surfaces S with boundary and $(g + \ell - 1)$ positive punctures $\{p_1, \dots, p_{g+\ell-1}\}$ and $(g + \ell - 1)$ negative punctures $\{q_1, \dots, q_{g+\ell-1}\}$ on ∂S, such that S is compact away from the punctures. We will consider J-holomorphic curves $u : S \to W_{\mathbb{D}}$ such that

(1) The map u is a smooth embedding;
(2) $u(\partial S) \subseteq C_\alpha \cup C_\beta$;

(3) u is nonconstant on each component of S;
(4) u has finite energy;
(5) for each i, $u^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ and $u^{-1}(\beta_i \times \{0\} \times \mathbb{R})$ consists of exactly one component of ∂S;
(6) $\lim_{w \to p_i} \pi_{\mathbb{R}} \circ u(w) = -\infty$ and $\lim_{w \to q_i} \pi_{\mathbb{R}} \circ u(w) = +\infty$.

Let $\overline{W}_{\mathbb{D}} = \Sigma \times \mathbb{D}$ denote the compactification of $W_{\mathbb{D}}$, and let $e_\alpha = \partial\mathbb{D} \cap \{\mathrm{Re}(z) \geq 0\}$ and $e_\beta = \partial\mathbb{D} \cap \{\mathrm{Re}(z) \leq 0\}$. We let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the homology classes of maps

$$\phi : (S, \partial S) \to (W_{\mathbb{D}}, C_\alpha \cup C_\beta)$$

which converge to $\mathbf{x}$ and $\mathbf{y}$ at the punctures, where two maps are considered homologous if they define the same element in

$$H_2(\overline{W}_{\mathbb{D}}, (\boldsymbol{\alpha} \times e_\alpha) \cup (\boldsymbol{\beta} \times e_\beta); \mathbb{Z}).$$

Given a $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we now define $\mathcal{M}(\phi)$ to be the moduli space of J-holomorphic curves which represent $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. The space $\mathcal{M}(\phi)$ has a free $\mathbb{R}$-action, so we denote by $\widehat{\mathcal{M}}(\phi)$ the space $\mathcal{M}(\phi)/\mathbb{R}$. We will often have to consider the compactification of spaces $\widehat{\mathcal{M}}(\phi)$, which we denote by $\overline{\widehat{\mathcal{M}}(\phi)}$.

There is a Maslov index $\mu : \pi_2(\mathbf{x}, \mathbf{y}) \to \mathbb{Z}$ which computes the expected dimension of the moduli space $\mathcal{M}(\phi)$. The Maslov formula is preserved by splicing disks together, i.e. if $\phi_1 \in \pi_2(\mathbf{x}, \mathbf{y})$ and $\phi_2 \in \pi_2(\mathbf{y}, \mathbf{z})$, then

$$\mu(\phi_1 * \phi_2) = \mu(\phi_1) + \mu(\phi_2).$$

See Section 4 of [Lip06] for technicalities on the index. For our purposes, it will be more useful to use combinatorial formulas for the index, which we will discuss in Section 4.4.

4.2. **The colored Heegaard Floer complexes.** Given a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we define the local multiplicity at a point $p \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ to be the algebraic intersection number

$$n_p(\phi) = \#(\phi^{-1}(\{p\} \times I \times \mathbb{R})).$$

Given an ℓ-tuple, $\mathbf{w}$, we define $n_{\mathbf{w}} : \pi_w(\mathbf{x}, \mathbf{y}) \to \mathbb{Z}^\ell$ by

$$n_{\mathbf{w}}(\phi) = (n_{w_1}(\phi), \dots, n_{w_\ell}(\phi)).$$

An essential feature of Heegaard Floer homology is that it is graded over Spin^c structures. Indeed the admissibility requirements are an unfortunate necessity as a single diagram can often not be chosen to compute all flavors of HF° for all Spin^c structures. There is a map from intersection points to Spin^c structures:

$$\mathfrak{s}_{\mathbf{w}} : \mathbb{T}_\alpha \cap \mathbb{T}_\beta \to \mathrm{Spin}^c(Y),$$

(cf. eg. [OS08, Section 3.3])

To each basepoint $w \in \mathbf{w}$, we associate the formal variable U_w. To achieve functoriality of the graph cobordism maps, we must consider colorings of the basepoints, i.e. Heegaard Floer complexes where we have formally set some of the variables equal to each other before taking homology. We also need to allow formal variables which don't correspond to any basepoints:

Definition 4.2. *A **coloring**, $(\sigma, \mathfrak{P})$ of a collection of basepoints $\mathbf{w}$ is a finite index set $\mathfrak{P}$ and a function $\sigma : \mathbf{w} \to \mathfrak{P}$. A **coloring** of a graph Γ is a map $\sigma : C(\Gamma) \to \mathfrak{P}$, where $C(\Gamma)$ denotes the set of connected components of Γ.*

We remark that a coloring $\sigma : \mathbf{w} \to \mathfrak{P}$ need not be surjective. When there is no chance of confusion, we will omit writing $\mathfrak{P}$ and just write σ for a coloring.

We will use the notation

$$\mathbb{Z}_2[U_{\mathbf{w}}] \stackrel{\mathrm{def}}{=} \mathbb{Z}_2[U_{w_1}, \dots, U_{w_\ell}]$$

and

$$\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}] \stackrel{\mathrm{def}}{=} \mathbb{Z}_2[U_{w_1}, \dots, U_{w_\ell}, U_{w_1}^{-1}, \dots, U_{w_\ell}^{-1}].$$

Similarly if $\mathfrak{P} = \{p_1, \dots, p_n\}$, we define

$$\mathbb{Z}_2[U_{\mathfrak{P}}] = \mathbb{Z}_2[U_{p_1}, \dots, U_{p_n}]$$

and $\mathbb{Z}_2[U_{\mathfrak{P}}, U_{\mathfrak{P}}^{-1}]$ similarly.

We make the following definition:

Definition 4.3. *Given a coloring $(\mathfrak{P},\sigma)$ of basepoints $\mathbf{w}$, we define $\mathcal{C}_{(\mathfrak{P},\sigma)}$ to be the $\mathbb{Z}_2[U_{\mathbf{w}},U_{\mathfrak{P}}]$ module defined by*

$$C_{\mathfrak{P},\sigma}=\mathbb{Z}_2[U_{\mathbf{w}},U_{\mathfrak{P}}]/I_{\mathfrak{P},\sigma}$$

where $I_{\mathfrak{P},\sigma}$ is the ideal generated by elements of the form $U_w-U_{\sigma(w)}$.

We can now define the Heegaard Floer complexes. Let $\mathcal{H}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w})$ be a strongly $\mathfrak{s}$-admissible diagram for $(Y,\mathbf{w},\mathfrak{s})$ (see Section 4.5 for a discussion of admissibility). We let $\widehat{CF}(\mathcal{H},\mathfrak{s})$ denote the free $\mathbb{Z}_2$-module generated by $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$. The coloring has no effect on the hat flavor.

For the other flavors, we first define the uncolored complexes. Define $CF^\infty(\mathcal{H},\mathfrak{s})$ to be the $\mathbb{Z}_2[U_{\mathbf{w}},U_{\mathbf{w}}^{-1}]$-module generated by intersection points $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, i.e. formal sums of elements of the form $\mathbf{x}\cdot U_{w_1}^{i_1}\cdots U_{w_\ell}^{i_\ell}$ where $i_1,\dots,i_\ell$ are arbitrary integers, and $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ is an intersection point with $s_{\mathbf{w}}(\mathbf{x})=\mathfrak{s}$. Define $CF^-(\mathcal{H},\mathfrak{s})$ to be the free $\mathbb{Z}[U_{\mathbf{w}}]$-submodule generated by intersection points $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, i.e. $CF^-(\mathcal{H},\mathfrak{s})$ consists of sums of elements of the form $\mathbf{x}\cdot U_{w_1}^{i_1}\cdots U_{w_\ell}^{i_\ell}$ where $i_1,\dots,i_\ell\geq 0$ and $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})=\mathfrak{s}$. We define $CF^+(\mathcal{H},\mathfrak{s})=CF^\infty(\mathcal{H},\mathfrak{s})/CF^-(\mathcal{H},\mathfrak{s})$.

We finally define the colored complexes

$$CF^\circ(\mathcal{H},\mathfrak{s},\mathfrak{P},\sigma)=\mathcal{C}_{\mathfrak{P},\sigma}\otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]}CF^\circ(\mathcal{H},\mathfrak{s}),$$

for $\circ=-,+,\infty$.

Let J_s be an almost complex structure on $\Sigma\times I\times\mathbb{R}$ which is appropriately generic. For $\widehat{CF}$, define

$$\widehat{\partial}_{\mathcal{H},\mathfrak{s},J_s}(\mathbf{x})=\sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1\\ n_{\mathbf{w}}(\phi)=\mathbf{0}}}\#\widehat{\mathcal{M}}(\phi)\cdot\mathbf{y}.$$

For the other flavors, define

$$\partial^\infty_{\mathcal{H},\mathfrak{s},J_s}(\mathbf{x})=\sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1\\ n_{\mathbf{w}}(\phi)=\mathbf{0}}}\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{y},$$

which induces differentials on CF^- and CF^+.

A Gromov compactness argument, (cf. [OS08, Lemma 4.3]) shows that $\partial^\circ_{\mathcal{H},\mathfrak{s},J_s}$ square to zero.

4.3. **Holomorphic triangle maps.** Given three sets of curves $\boldsymbol{\alpha},\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ inside of a surface Σ, one can consider maps which count holomorphic triangles inside of $\mathrm{Sym}^{g+\ell-1}(\Sigma)$ with boundary on $\mathbb{T}_\alpha,\mathbb{T}_\beta$ and $\mathbb{T}_\gamma$. Following [Lip06], one can also compute these maps in the cylindrical formulation, which we briefly describe in this section.

A **multipointed Heegaard triple**, $\mathcal{T}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{w})$, is a collection where each $\boldsymbol{\alpha},\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are collections of $(g+\ell-1)$ pairwise disjoint, simple closed curves which each span a g-dimensional sublattice of $H_1(\Sigma;\mathbb{Z})$. We also require each σ_i and τ_j to be transverse for $\sigma,\tau\in\{\alpha,\beta,\gamma\}$. We also will need to assume an admissibility assumption, which we discuss in Section 4.5.

The triple $\mathcal{T}$ specifies three Heegaard diagrams, which we denote by $\mathcal{T}_{\alpha\beta}$, $\mathcal{T}_{\alpha\gamma}$ and $\mathcal{T}_{\beta\gamma}$. By attaching handles along these curves and then additional 0- and 3-handles, the triple specifies 3 closed manifolds $Y_{\alpha\beta}$, $Y_{\alpha\gamma}$, and $Y_{\beta\gamma}$.

The triple also specifies a compact 4-manifold $X_{\alpha\beta\gamma}$ with boundary

$$Y_{\alpha\beta}\sqcup Y_{\beta\gamma}\sqcup -Y_{\alpha\gamma}$$

obtained as follows: let U_α,U_β, U_γ denote the handlebodies specified by attaching handles along the α,β or γ curves respectively. Let $\overline{\Delta}$ be the closed triangle with edges e_α,e_β and e_γ. By gluing $U_\alpha\times[0,1],U_\beta\times[0,1]$ and $U_\gamma\times[0,1]$ to $\Sigma\times\overline{\Delta}$, we obtain $X_{\alpha\beta\gamma}$.

We let $v_{\alpha\beta},v_{\beta\gamma}$ and $v_{\alpha\gamma}$ denote the vertices of $\overline{\Delta}$ and we let W_Δ denote $\Sigma\times\Delta$. Given intersection points $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta,\mathbf{y}\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma$, and $\mathbf{z}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma$, we let $\pi_2(\mathbf{x},\mathbf{y},\mathbf{z})$ denote homology classes of maps $\psi:(S,\partial S)\to(\Sigma\times\Delta,\boldsymbol{\alpha}\times e_\alpha\cup\boldsymbol{\beta}\times e_\beta\cup\boldsymbol{\gamma}\cup e_\gamma)$ where S is a Riemann surface with $3(g+\ell-1)$ punctures on the boundary, labelled as $p_i^{\alpha\beta},p_i^{\beta\gamma}$, and $p_i^{\alpha\gamma}$. We further require that ψ extends continuously to the compactification $\overline{\psi}:(\overline{S},\overline{\partial S})\to\Sigma\times\overline{\Delta}$ such that $\pi_{\overline{\Delta}}\overline{\psi}(p_i^{\alpha\beta})=v_{\alpha\beta}$ and similarly for the $p_i^{\alpha\gamma}$ and $p_i^{\beta\gamma}$ points.

To compute the triangle maps in the cylindrical setting, we wish to count holomorphic representatives of a homology class $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$. We let $\mathcal{M}(\psi)$ denote the set of holomorphic curves $u : (S, \partial S) \to W_\Delta$ for an appropriately generic almost complex structure J on W_Δ satisfying the following:

(1) u is homologous to ψ after forgetting the almost complex structure;
(2) u is an embedding;
(3) u has finite energy;
(4) for each $\sigma \in \{\alpha, \beta, \gamma\}$ and $i = 1, \dots, (g+\ell-1)$, the set $u^{-1}(\sigma_i \times e_\sigma)$ consists of exactly one component of ∂S;
(5) u is nonconstant on each component of S.

In general, the triangle maps in Heegaard Floer homology must be defined with reference to a particular 4-dimensional Spin^c structure. Following [OS06, Sec. 2.2], or [Lip06, Sec. 10.1.2], there is a map

$$s_{\mathbf{w}} : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathrm{Spin}^c(X_{\alpha\beta\gamma}),$$

such that

$$s_{\mathbf{w}}(\psi)|_{Y_{\alpha\beta}} = s_{\mathbf{w}}(\mathbf{x})$$

and similarly for $\mathbf{y}$ and $\mathbf{z}$.

Using this, we can thus define triangle maps

$$\begin{aligned}
&\widehat{f}_{\alpha\beta\gamma,\mathfrak{s}} : \widehat{CF}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}|_{\alpha\beta}) \otimes_{\mathbb{Z}_2} \widehat{CF}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to \widehat{CF}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
&f^{\infty}_{\alpha\beta\gamma,\mathfrak{s}} : CF^{\infty}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}|_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}]} CF^{\infty}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{\infty}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
&f^{-}_{\alpha\beta\gamma,\mathfrak{s}} : CF^{-}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]} CF^{-}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{-}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
&f^{+-}_{\alpha\beta\gamma,\mathfrak{s}} : CF^{+}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]} CF^{-}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{+}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
&f^{-+}_{\alpha\beta\gamma,\mathfrak{s}} : CF^{-}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]} CF^{+}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{+}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),
\end{aligned}$$

which we define using the formulas

$$\widehat{f}_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x} \otimes \mathbf{y}) = \sum_{\mathbf{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\substack{\psi \in \pi_2(\mathbf{x},\mathbf{y},\mathbf{z}) \\ \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s} \\ \mu(\psi) = 0 \\ n_{\mathbf{w}}(\psi) = \mathbf{0}}} \#\mathcal{M}(\psi) \cdot \mathbf{z}$$

and

$$f^{\circ}_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x} \otimes \mathbf{y}) = \sum_{\mathbf{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\substack{\psi \in \pi_2(\mathbf{x},\mathbf{y},\mathbf{z}) \\ \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s} \\ \mu(\psi) = 0}} \#\mathcal{M}(\psi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)} \cdot \mathbf{z}$$

for $\circ = +, -, \infty$.

Taking homology, we get a map

$$\widehat{F}_{\alpha\beta\gamma,\mathfrak{s}} : \widehat{HF}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2} \widehat{HF}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to \widehat{HF}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma})$$

as well as maps $F^{\infty}_{\alpha\beta\gamma,\mathfrak{s}}, F^{-}_{\alpha\beta\gamma,\mathfrak{s}}, F^{+-}_{\alpha\beta\gamma,\mathfrak{s}}$ and $F^{-+}_{\alpha\beta\gamma,\mathfrak{s}}$ defined analogously.

4.4. **Acceptable domains and combinatorial Maslov indices.** One of the key features of Heegaard Floer homology is that homotopy classes of disks and triangles, as well as their Maslov indices can be computed completely combinatorially.

4.4.1. *Domains.* Let $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ have connected components $r_1, \dots, r_n$ and let R_i be the closure of r_i. We call R_i **regions**. A **domain** is a formal $\mathbb{Z}$-linear combination of the R_i. To a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we can associate a domain

$$\mathcal{D}(\phi) = \sum_{i=1}^{n} n_{p_i}(\phi) R_i,$$

where p_i is any point in the region R_i. The domain of a triangle $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is defined analogously.

Given an intersection point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and an intersection point $x \in \alpha_i \cap \beta_j$, we define

$$\delta(\mathbf{x}, x) = \begin{cases} 1 & \text{if } x \in \mathbf{x}, \\ 0 & \text{otherwise} \end{cases}.$$

Suppose $\mathcal{D}$ is a domain, and $x \in \alpha_i \cap \beta_j$ is an intersection point. There are four regions regions with a corner at x. While we traverse counterclockwise a small circle centered at x, in two of these regions we start at α_i and end at β_j. Let $m^{\alpha\beta}(\mathcal{D}, x)$ be the sum of the multiplicities in these regions. Similarly in the other two regions, we start at β_j and end at α_i. We let $m^{\beta\alpha}(\mathcal{D}, x)$ denote the sum of the multiplicities in these regions.

We say that a domain $\mathcal{D}$ is **acceptable for** $(\mathbf{x}, \mathbf{y})$ if

$$m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) = \delta(\mathbf{x}, x) - \delta(\mathbf{y}, x) \tag{1}$$

for all i, and j and all $x \in \alpha_i \cap \beta_j$.

Lemma 4.4. *There is a bijection between domains $\mathcal{D}$ which are acceptable for $(\mathbf{x}, \mathbf{y})$ and elements of $\pi_2(\mathbf{x}, \mathbf{y})$.*

There is a similar relationship between acceptable domains and triangles.

Lemma 4.5. *There is a bijection between classes in $\pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ and $\mathcal{D}$ which satisfy*

$$\begin{aligned} m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) &= \delta(\mathbf{x}, x), \\ m^{\beta\gamma}(\mathcal{D}, y) - m^{\gamma\beta}(\mathcal{D}, y) &= \delta(\mathbf{y}, y), \\ m^{\gamma\alpha}(\mathcal{D}, z) - m^{\alpha\gamma}(\mathcal{D}, z) &= \delta(\mathbf{z}, z) \end{aligned}$$

for all intersection points x, y, and z.

Definition 4.6. *We say that a homology disk $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ is a **periodic class** if $n_{\mathbf{w}}(\phi) = \mathbf{0}$. We say a domain $\mathcal{D}$ is a **periodic domain** if $m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) = 0$ for all $x \in \alpha_i \cap \beta_j$ and $n_{\mathbf{w}}(\phi) = \mathbf{0}$.*

4.4.2. *Combinatorial Maslov indices.* The Maslov index $\mu : \pi_2(\mathbf{x}, \mathbf{y}) \to \mathbb{Z}$ and $\mu : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathbb{Z}$ compute the expected dimension of the moduli spaces of holomorphic triangles. The definition of μ, (cf. eg. [OS04b], [Lip06]) involves taking the index of the pseudodifferential $\overline{\partial}$ operator, which normally is not computationally tractable, but for our purposes it is sufficient to know that there is a combinatorial computation of the Maslov index, due in part to Lipshitz in [Lip06] and Sarkar in [Sar11a].

One first defines the Euler measure of a region R by

$$e(R) = \chi(R) - ac(R)/4 + ob(R)/4$$

where $ac(R)$ denotes the number of acute angles of R, and $ob(R)$ denotes the number of obtuse corners of R. We then extend e linearly to all domains.

A first formula for the Maslov index of a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is due to [Lip06, Cor. 4.3]:

$$\mu(\phi) = e(\mathcal{D}) + \sum_{x \in \mathbf{x}} n_x(\mathcal{D}) + \sum_{y \in \mathbf{y}} n_y(\mathcal{D}), \tag{2}$$

where $\mathcal{D} = \mathcal{D}(\phi)$ and $n_p(\mathcal{D})$ denotes the average of the multiplicities in the four regions with vertex $p \in \alpha_i \cap \beta_j$.

Another useful formula which we will need holds for certain classes of homology classes. Let $\pi_2^{\alpha}(\mathbf{x})$ denote homology classes of disks with boundary in $\boldsymbol{\alpha}$ and define $\pi_2^{\beta}(\mathbf{x})$ similarly. For $\phi \in \pi_2^{\beta}(\mathbf{x})$ or $\phi \in \pi_2^{\alpha}(\mathbf{x})$, following [OS08, Lem. 5.4], we have

$$\mu(\phi) = 2 \sum_{i=1}^{\ell} n_{w_i}(\phi). \tag{3}$$

4.5. **Admissibility.** In order to compute HF°, one must use diagrams with certain admissibility properties. To compute HF^+ and $\widehat{HF}$ one needs a weaker form of admissibility, usually referred to as *weak admissibility* or just *admissibility*. To compute HF^- and HF^∞, one needs to use a stronger form of admissibility, usually referred to as *strong $\mathfrak{s}$-admissibility*, which is specified to a particular Spin^c structure $\mathfrak{s}$. To compute cobordism maps in our framework, it will be sufficient to use weak admissibility for $\widehat{HF}$, but we will need strong admissibility for HF^+, HF^- and HF^∞. The reason we can't use weak admissibility for HF^+ is essentially due to technical neck stretching arguments used to define the 1- and 3-handle maps.

4.5.1. *Admissibility for Heegaard diagrams.* Before we state the admissibility condition, we must first recall the Chern class formula for the Maslov index proven by Ozsváth and Szabó.

In their analysis of topological disks, Ozsváth and Szabó produce the following exact sequence [OS08, Eq. 4]:

$$0 \to \mathbb{Z} \to \pi_2^\alpha(\mathbf{x}) \oplus \pi_2^\beta(\mathbf{x}) \to \pi_2(\mathbf{x}, \mathbf{x}) \to H_2(Y; \mathbb{Z}) \to 0.$$

Given a $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$, we let $H(\phi)$ denote its image in $H_2(Y; \mathbb{Z})$ under the above map. The following Lemma is essentially [OS04a, Prop. 7.5], except for multipointed diagrams:

Lemma 4.7. *Suppose that $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$. Then*

$$\mu(\phi) = \langle c_1(\mathfrak{s}), H(\phi) \rangle + 2 \sum_{i=1}^{\ell} n_{w_i}(\phi).$$

Proof. This follows analogously to the proof of the case of a single basepoint ([OS04a, Prop. 7.5]). More precisely, instead of equation (5) in [OS04a], we replace $n_z(\mathcal{P} + m[\Sigma])$ with $\sum_{i=1}^{\ell} n_{w_i}(\mathcal{P} + m[\Sigma])$, since we must remove balls centered around the trajectories from index 0 to index 3 critical points passing through each basepoint w_i.

Performing the analogous computation as in [OS04a, Prop. 7.5] to produce their equation (5), we thus have if $\mathcal{D} = \mathcal{D}(\phi)$ is the domain of a $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ (not necessarily periodic), then

$$\langle c_1(\mathfrak{s}), H(\phi) \rangle = e(\mathcal{D} + m[\Sigma]) + 2 \left(\sum_{x \in \mathbf{x}} n_x(\mathcal{D} + m[\Sigma]) \right) - 2 \left(\sum_{i=1}^{\ell} n_{w_i}(\mathcal{D} + m[\Sigma]) \right)$$

for sufficiently large m. By expanding this and using Equations (2) and (3), we recover the desired formula. □

The strong notion of admissibility we will use for Heegaard diagrams is the following:

Definition 4.8. *A diagram $\mathcal{H}$ is **strongly $\mathfrak{s}$-admissible** for a* Spin^c *structure $\mathfrak{s}$ if for every nontrivial periodic domain $\mathcal{P}$ with*

$$\langle c_1(\mathfrak{s}), H(\mathcal{P}) \rangle = 2N > 0,$$

then $\mathcal{P}$ has some local multiplicity strictly greater than N.

The following is only a small modification of [OS04b, Lem. 4.14]:

Lemma 4.9. *If $\mathcal{H}$ is strongly admissible for $\mathfrak{s}$, then for a fixed j there are only finitely many disks $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = j$ with $\mathcal{D}(\phi) \geq 0$.*

Proof. Fix any $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$. Any other disk $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$ satisfies

$$\mathcal{D}(\phi) = \mathcal{D}(\psi) + \mathcal{P} + \mathcal{A}$$

for a periodic domain $\mathcal{P}$ and a domain $\mathcal{A}$ which is a $\mathbb{Z}$-linear combination of the components of $\Sigma \setminus \boldsymbol{\alpha}$. If we assume $\mathcal{D}(\phi) \geq 0$, then we have

$$-\mathcal{D}(\psi) \leq \mathcal{P} + \mathcal{A}.$$

Suppose we have an infinite sequence of distinct ϕ_n. This yields an infinite sequence of distinct pairs $(\mathcal{P}_i, \mathcal{A}_i)$ satisfying the above inequality.

An infinite sequence of distinct such domains $\mathcal{P}_i + \mathcal{A}_i$ would then need to satisfy $\|\mathcal{P}_i + \mathcal{A}_i\| \to \infty$ where $\| \cdot \|$ denotes the Euclidean 2-norm, since they are elements of a finitely generated integer lattice. By

renormalizing, note that both $\mathcal{P}_i+\mathcal{A}_i$ and $\mathcal{A}_i$ are bounded below, so we can extract a convergent subsequences so that $\mathcal{P}_i/\|\mathcal{P}_i+\mathcal{A}_i\|$ and $\mathcal{A}_i/\|\mathcal{P}_i+\mathcal{A}_i\|$ both converge to real domains $\mathcal{P}_\infty$ and $\mathcal{A}_\infty$ respectively. Since

$$0=\mu(\mathcal{P}_i+\mathcal{A}_i)=\mu(\mathcal{A}_i)+\langle c_1(\mathfrak{s}),H(\mathcal{P}_i)\rangle=\left(\sum_{j=1}^{\ell}2n_{w_j}(\mathcal{A}_i)\right)+\langle c_1(\mathfrak{s}),H(\mathcal{P}_i)\rangle$$

and the above equation is preserved by normalization, we know that

$$\left(\sum_{j=1}^{\ell}2n_{w_j}(\mathcal{A}_\infty)\right)+\langle c_1(\mathfrak{s}),H(\mathcal{P}_\infty)\rangle=0.$$

Since $\mathcal{A}_i$ were bounded below, we know that $\mathcal{A}_\infty\geq 0$. It is easy to see that given a real valued such pair $(\mathcal{A}_\infty,\mathcal{P}_\infty)$, there is a rational one $(\mathcal{A},\mathcal{P})$ also with $\mathcal{A}\geq 0$ and $\mathcal{A}+\mathcal{P}\geq 0$. By clearing denominators we can assume that $\mathcal{A}$ and $\mathcal{P}$ are integral. We have, however, that

$$\langle c_1(\mathfrak{s}),H(-\mathcal{P})\rangle=2N$$

where $N=\sum_{j=1}^{\ell}n_{w_j}(\mathcal{A})\geq 0$. Hence $-\mathcal{P}$ has a multiplicity which is greater than N by strong admissibility. Since $\mathcal{A}\geq 0$, this contradicts the fact that $\mathcal{A}+\mathcal{P}\geq 0$. □

4.5.2. *Admissibility for higher multidiagrams.* Ozsváth and Szabó produced similar admissibility conditions for Heegaard triples (cf. [OS04b, Def. 8.8]), and Heegaard quadruples (cf. [OS04b, Sec. 8.4.2.]). Instead of stating these directly, we will state a slightly more general description of higher Heegaard multidiagrams.

Definition 4.10. *An abstract* ***Heegaard multi-diagram*** $(\Sigma,\boldsymbol{\eta}_0,\dots,\boldsymbol{\eta}_n,\mathbf{w})$ *is a surface* Σ *with sets of curves* $\boldsymbol{\eta}_i=\{\eta_i^1,\dots,\eta_i^{g+\ell-1}\}$ *such that*

(1) Σ *is a closed, oriented, genus* g *surface with* $\mathbf{w}\subseteq\Sigma$*;*
(2) $\boldsymbol{\eta}_i=(\eta_i^1,\dots,\eta_i^{g+\ell-1})$ *is a collection of disjoint, simple closed curves on* Σ*;*
(3) no more than two curves from any collection intersect at a single point;
(4) all curves are transverse to each other.

Given a Heegaard multi-diagram $\mathcal{M}=(\Sigma,\boldsymbol{\eta}_0,\dots,\boldsymbol{\eta}_n,\mathbf{w})$, we can construct a compact 4-manifold $X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n}$ by gluing copies of $U_{\boldsymbol{\eta}_i}\times e_{\boldsymbol{\eta}_i}$ to $G_n\times\Sigma$, where G_n is the regular $(n+1)$-gon, with sides labelled $e_{\boldsymbol{\eta}_i}$, counterclockwise. The manifold $X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n}$ has boundary equal to

$$Y_{\boldsymbol{\eta}_0\boldsymbol{\eta}_1}\sqcup\dots\sqcup Y_{\boldsymbol{\eta}_{n-1}\boldsymbol{\eta}_n}\sqcup Y_{\boldsymbol{\eta}_n\boldsymbol{\eta}_0}.$$

Similar to the map from intersection points on a Heegaard diagram to three dimensional Spin^c structures, there is a map

$$\mathfrak{s}_{\mathbf{w}}:\pi_2(\mathbf{x}_0,\dots,\mathbf{x}_n)\to\mathrm{Spin}^c(X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n})$$

such that

$$\mathfrak{s}_{\mathbf{w}}(\phi)|_{Y_{\boldsymbol{\eta}_j\boldsymbol{\eta}_{j+1}}}=\mathfrak{s}_{\mathbf{w}}(\mathbf{x}_j).$$

Such a map is defined for triangles and rectangles in [OS04b] (and those are the only cases that we will need), and for higher polygons in [GW08]. The following is analogous to [GW08, Proposition 3.9].

Lemma 4.11. *Let* $\psi\in\pi_2(\mathbf{x}_0,\dots,\mathbf{x}_n)$ *and* $\psi'\in\pi_2(\mathbf{x}'_0,\dots,\mathbf{x}'_n)$ *be two topological* $(n+1)$*-gons. Then* $\mathfrak{s}(\psi)=\mathfrak{s}(\psi')$ *iff*

$$\psi'=\psi+\phi_0+\cdots\phi_n$$

for classes $\phi_i\in\pi_2(\mathbf{x}_i,\mathbf{x}'_i)$.

Proof. This follows as in [OS04b, Prop. 8.5] for triangles, or [GW08, Prop. 3.9] for higher n-gons. □

For 3-gons and 4-gons, Ozsváth and Szabó define notions of strong admissibility. In more generality, following [GW08] we can generalize these notions to higher n-gons. To be able to use counts of holomorphic n-gons to define maps for $n\geq 4$, we need a strong admissibility condition not just for a single Spin^c structure, but instead for certain families of Spin^c structures.

Given an n-gon G, we can consider a collection of disjoint arcs A on G such that each arc intersects exactly two edges of G, no two arcs intersect each other, and such that each component of $G\setminus A$ intersects exactly

three sides of ∂G. We will call such a collection a **complete set of dividing arcs on** G. Given a complete set of dividing arcs, we get a decomposition of $X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n}$ into triangular cobordisms of the form $X_{\boldsymbol{\eta}_i\boldsymbol{\eta}_j\boldsymbol{\eta}_k}$. We call such a decomposition of $X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n}$ a **complete decomposition into triangular cobordisms**. An example of a decomposition for the cobordism $X_{\alpha\beta\gamma\delta}$ is shown in Figure 2.

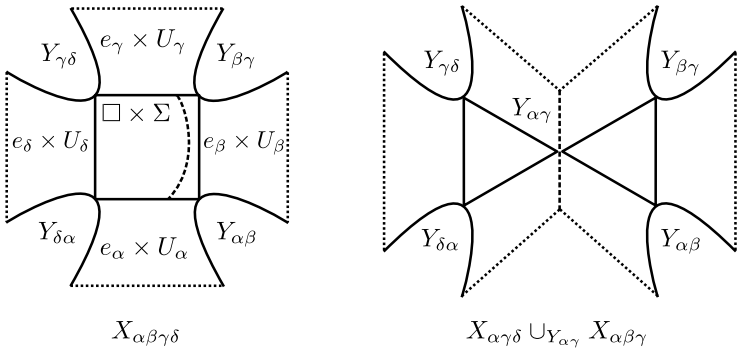

FIGURE 2. Decomposing the cobordism $X_{\alpha\beta\gamma\delta}$ into two triangular cobordisms.

Definition 4.12. *We define two* Spin^c *structures* $\mathfrak{s}, \mathfrak{s}' \in \text{Spin}^c(X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n})$ *to be* ***equivalent*** *if for every decomposition of* $X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n}$ *into triangular cobordisms, we have*

$$\mathfrak{s}|_{X_\Delta} = \mathfrak{s}'|_{X_\Delta}$$

for each of the triangular cobordisms X_Δ*. We use the letter* $\mathfrak{S}$ *to denote such an equivalence class of* Spin^c *structures.*

Remark 4.13. *Given a rectangular cobordism* $X_{\alpha\beta\gamma\delta}$*, there are exactly two complete decompositions of* $X_{\alpha\beta\gamma\delta}$*, one of which is given by the arc from* e_γ *to* e_α*, and the other is given by the arc from* e_δ *to* e_β*. Using the Meyer-Vietoris sequence, we thus see that equivalence classes* $\mathfrak{S}$ *in* Spin$^c(X_{\alpha\beta\gamma\delta})$ *correspond exactly to* $\delta H^1(Y_{\alpha\gamma}) + \delta H^1(Y_{\beta\delta}))$ *orbits within* Spin$^c(X_{\alpha\beta\gamma\delta})$*, so our definition coincides for rectangles with the definition given in* [OS04b]*. Similarly note that there is only one triangular decomposition of a triangle cobordism, and hence two* Spin^c *structures are equivalent on a triangular cobordism iff they are equal.*

Definition 4.14. *We say that a multi-diagram* $\mathcal{M} = (\Sigma, \boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_n, \mathbf{w})$ *is strongly* $\mathfrak{S}$*-admissible for an equivalence class* $\mathfrak{S}$ *of* Spin^c *structures* $\mathfrak{s} \in \text{Spin}^c(X_{\boldsymbol{\eta}_0\dots\boldsymbol{\eta}_n})$ *if for each* $\mathfrak{s} \in \mathfrak{S}$*, if* $\mathcal{P}$ *is a domain which can be written as a sum of doubly-periodic domains*

$$\mathcal{P} = \sum_{\{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}\} \subseteq \{\boldsymbol{\eta}_0,\dots,\boldsymbol{\eta}_n\}} \mathcal{P}_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}},$$

with

$$\sum_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}} \langle c_1(\mathfrak{s}|_{Y_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}}}), H(\mathcal{P}_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}})\rangle = 2N \geq 0,$$

then $\mathcal{P}$ *has some local multiplicity strictly greater than* N*.*

As remarked in [OS04b], in the case of a Heegaard quadruple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{w})$, a $\delta H^1(Y_{\alpha\gamma}) + \delta H^1(Y_{\beta\delta})$ orbit may consist of infinitely many Spin^c structures, so we cannot necessarily perform the standard winding trick to achieve strong $\mathfrak{S}$-admissibility. For most of the Heegaard quadruples we will consider, one can achieve strong admissibility by winding. The following is explained in [OS04b, Sec. 8.4.2]:

Remark 4.15. *If we assume that* $\delta H^1(Y_{\beta\delta})|_{Y_{\alpha\gamma}} = 0$ *and* $\delta H^1(Y_{\alpha\gamma})|_{Y_{\beta\delta}} = 0$*, then by winding sufficiently, we can achieve strong* $\mathfrak{S}$*-admissibility of a Heegaard quadruple.*

The following is analogous to [OS04b, Lem. 8.9]:

Lemma 4.16. *Suppose that* $\mathcal{M} = (\Sigma, \boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_n, \mathbf{w})$ *is a strongly* $\mathfrak{s}$*-admissible Heegaard multidiagram and suppose that* j *is fixed. Then there are finitely many homotopy classes of* $(n+1)$*-gons* $\psi \in \pi_2(\mathbf{x}_0, \dots, \mathbf{x}_n)$ *with* $\mu(\psi) = j, \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$ *and* $\mathcal{D}(\psi) \geq 0$*.*

Proof. The proof is analogous to [OS04b, Lem. 8.9] in the case of triangles. For higher $(n+1)$-gons, the proof is analogous. Fix a Spin^c structure on $X_{\boldsymbol{\eta}_0,\dots,\boldsymbol{\eta}_n}$, an integer j and a reference $(n+1)$-gon ψ_0 with $\mathfrak{s}_{\mathbf{w}}(\psi_0) = \mathfrak{s}$. By Lemma 4.11 if ψ is any other $(n+1)$-gon with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, then we can write

$$\psi = \psi_0 + \phi_0 + \cdots + \phi_n$$

with $\phi_i \in \pi_2(\mathbf{x}_i, \mathbf{x}_i)$. Note that since each $\phi_i \in \pi_2(\mathbf{x}_i, \mathbf{x}_i)$, as domains we have that each $\phi_i \in \pi_2(\mathbf{y}, \mathbf{y})$ for any $\mathbf{y}$, because the vertex relations in Equation 1 are satisfied. Hence in particular we can write

$$\psi = \psi_0 + \mathcal{A} + \mathcal{P}$$

where $\mathcal{A}$ is an α-boundary degeneration and $\mathcal{P}$ is a periodic domain. The proof then proceeds identically to the proof of Lemma 4.9.

□

4.6. **Heegaard moves and naturality of the 3-manifold invariants.** Given a multi-based 3-manifold $(Y, \mathbf{w})$, there are many different Heegaard diagrams for $(Y, \mathbf{w})$. The following is standard:

Proposition 4.17. *Any two embedded Heegaard diagrams $(\Sigma_1, \boldsymbol{\alpha}_1, \boldsymbol{\beta}_1, \mathbf{w})$ and $(\Sigma_2, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \mathbf{w})$ can be connected by a sequence of the following moves:*

1. *isotopies of $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves not passing over the $\mathbf{w}$ basepoints;*
2. *handleslides of the $\boldsymbol{\alpha}$ curves across each other or handleslides of the $\boldsymbol{\beta}$ curves across each other;*
3. *(1,2)-stabilizations or destabilizations;*
4. *isotopies of $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ within Y which are fixed on $\mathbf{w}$.*

We use the phrase "$(1,2)$-stabilization" to distinguish from free-stabiliation, which we will sometimes call $(0,3)$-stabilization, which occurs when we add a basepoint.

In it's original formulation in [OS04b], to an appropriately generic almost complex structure on $\mathrm{Sym}^k(\Sigma)$, Ozsváth and Szabó defined maps corresponding to moves (1)-(3) as well maps associated to changing the almost complex structure. Given two Heegaard diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$, as well as suitably generic almost complex structures J_i on $\mathrm{Sym}^{g_i}(\Sigma_i)$, Ozsváth and Szabó showed that there exists a chain homotopy equivalence

$$\Phi_{\mathcal{H}_1 \to \mathcal{H}_2} : CF^\circ_{J_1}(\mathcal{H}_1, \mathfrak{s}) \to CF^\circ_{J_2}(\mathcal{H}_2, \mathfrak{s}),$$

which is a composition of maps corresponding to moves of the above form, though they didn't show that the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ was independent of the choice of intermediate diagrams. This was resolved in [JT12] by Juhász and Thurston:

Theorem 4.18 ([JT12, Thm. 2.39])**.** *The map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ is independent up to chain homotopy of the choice of intermediate diagrams and almost complex structures.*

We will now describe the maps in more detail. Given sets of attaching curves $\boldsymbol{\nu}$ on Σ, we write $\boldsymbol{\nu} \sim \boldsymbol{\nu}'$ if $\boldsymbol{\nu}$ and $\boldsymbol{\nu}'$ are related by a sequence of handleslides or isotopies which don't cross $\mathbf{w}$. Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple with $\boldsymbol{\beta} \sim \boldsymbol{\gamma}$. Note that $X_{\alpha\beta\gamma}$ has a subset of the boundary diffeomorphic to $Y_{\beta\gamma} \simeq \sqcup_i (S^1 \times S^2)^{\#n}$, which can be "filled in" with some number of 3-handles and 4-handles. There is a bijection from Spin^c structures on $Y_{\alpha\beta}$ and Spin^c structures on $X_{\alpha\beta\gamma}$ which extend over the additional 3-handles and 4-handles. Given $\mathfrak{s} \in \mathrm{Spin}^c(Y_{\alpha\beta})$ we let $\overline{\mathfrak{s}}$ denote the unique such extension on $X_{\alpha\beta\gamma}$. A similar comment holds if $\boldsymbol{\alpha} \sim \boldsymbol{\beta}$, and we use the same notation for the extension of a Spin^c structure $\mathfrak{s}$ on $Y_{\beta\gamma}$ to $X_{\alpha\beta\gamma}$.

The following lemma is proven by Juhász and Thurston [JT12, Lem. 9.2.]:

Lemma 4.19. *Let $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ be a strongly (resp. weakly) $\mathfrak{s}$-admissible Heegaard triple. If $\boldsymbol{\beta} \sim \boldsymbol{\gamma}$ then $HF^-(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_0)$ and $\widehat{HF}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_0)$ admit "top" degree generators $\Theta_{\beta\gamma}$ for $\circ = \widehat{\ }, \infty, -$, and we write $\Psi^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\gamma}}$ for the map*

$$\Psi^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\gamma}} = F^\circ_{\alpha\beta\gamma,\overline{\mathfrak{s}}}(\cdot \otimes \Theta_{\beta\gamma}) : HF^\circ(\mathcal{T}_{\alpha\beta}, \mathfrak{s}) \to HF^\circ(\mathcal{T}_{\alpha\gamma}, \overline{\mathfrak{s}}|_{\alpha\gamma}).$$

Similarly if $\boldsymbol{\alpha} \sim \boldsymbol{\beta}$, we define

$$\Psi^{\boldsymbol{\alpha}\to\boldsymbol{\beta}}_{\boldsymbol{\gamma}} = F^\circ_{\alpha\beta\gamma,\overline{\mathfrak{s}}}(\Theta_{\alpha\beta} \otimes \cdot) : HF^\circ(\mathcal{T}_{\beta\gamma}, \mathfrak{s}) \to HF^\circ(\mathcal{T}_{\alpha\gamma}, \overline{\mathfrak{s}}_{\alpha\gamma}).$$

Juhász and Thurston also reduced the number of types of Heegaard moves which one needs to perform to go between different Heegaard diagrams. Originally Ozsváth and Szabó defined continuation maps $\Gamma^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\gamma}}$ in [OS04b] to be the transition maps for move (1), above. These involved counting pseudoholomorphic disks with dynamic boundary conditions. According to [JT12, Lem. 9.5], if the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w})$ achieves admissibility, then

$$\Gamma^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'} = \Psi^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'}.$$

If we don't assume admissibility, the map $\Gamma^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'}$ is instead equal to a composition of triangle maps over Heegaard triples which do achieve admissiblity (cf. [JT12, Prop. 9.9]).

More generally, given admissible diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ and $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w})$ with $\boldsymbol{\alpha} \sim \boldsymbol{\alpha}'$ and $\boldsymbol{\beta} \sim \boldsymbol{\beta}'$, Juhász and Thurston define a map

$$\Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'} : HF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s}) \to HF^\circ(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}, \mathfrak{s})$$

which is a composition of triangle maps, generalizing the maps defined for individual moves. In particular this allows us to avoid using the continuation maps.

We now provide notation for the remaining maps. If $\mathcal{H}_2$ is a stabilization of $\mathcal{H}_1$, then we let $\sigma : HF^\circ(\mathcal{H}_1, \mathfrak{s}) \to HF^\circ(\mathcal{H}_2, \mathfrak{s})$ denote the stabilization isomorphism. If ϕ_t is an isotopy of Y fixing $\mathbf{w}$, and $\phi = \phi_1$, we define $\phi_* : HF^\circ(\mathcal{H}, \mathfrak{s}) \to HF^\circ(\phi_*(\mathcal{H}), \mathfrak{s})$ by the tautological formula

$$\phi_*(x_1, \dots, x_{g+\ell-1}) = (\phi_* x_1, \dots, \phi_* x_{g+\ell-1})$$

for an intersection point $(x_1, \dots, x_{g+\ell-1})$.

We summarize these comments into the following lemma:

Lemma 4.20. *Suppose that $\mathcal{H}_1$ and $\mathcal{H}_2$ are two embedded Heegaard diagrams for $(Y, \mathbf{w})$ and J_1 and J_2 are two suitably generic almost complex structures on $W_{\mathbb{D}}$. Then the map $\Phi_{\mathcal{H}_1\to\mathcal{H}_2} : CF^\circ_{J_1}(\mathcal{H}_1, \mathfrak{s}) \to CF^\circ_{J_2}(\mathcal{H}_2, \mathfrak{s})$ can be written as a composition of the following maps:*

1. *change of almost complex structure maps $\Phi_{J_s\to J'_s}$;*
2. *triangle maps of the form $F_{\alpha\beta\gamma,\mathfrak{s}}(\Theta_{\alpha\beta} \otimes \cdot)$ if $\boldsymbol{\alpha} \sim \boldsymbol{\beta}$ or $F_{\alpha\beta\gamma,\mathfrak{s}}(\cdot \otimes \Theta_{\beta\gamma})$ if $\boldsymbol{\beta} \sim \boldsymbol{\gamma}$;*
3. *stabilization maps σ or their inverses σ^{-1};*
4. *diffeomorphism maps ϕ_* for diffeomorphisms $\phi = \phi_1$ where ϕ_t is an isotopy of $(Y, \mathbf{w})$ fixing $\mathbf{w}$ for all t.*

5. Homological actions on the Heegaard Floer groups

We now introduce maps on $HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{s})$ which we call the relative homology maps. These roughly correspond to an action of $H_1(Y/\sigma; \mathbb{Z})$ on $HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{s})$ where σ is a coloring of the basepoints $\mathbf{w}$ and by Y/σ we mean the topological space given by identifying basepoints of the same color in Y. Given a path λ in Y between w_1 and w_2 which is immersed in the underlying surface of a Heegaard diagram $\mathcal{H}$ for Y, we define a map $A_\lambda : CF^\circ(\mathcal{H}, \mathfrak{s}) \to CF^\circ(\mathcal{H}, \mathfrak{s})$ which satisfies

$$A_\lambda \partial + \partial A_\lambda = U_{w_1} + U_{w_2}.$$

For complexes colored so that $U_{w_1} = U_{w_2}$, we get that A_λ is a chain map. For a closed loop γ inside of the underlying surface of $\mathcal{H}$, we get the familiar map A_γ, which is a chain map satisfying $A_\gamma \circ A_\gamma \simeq 0$.

These maps generalize the more familiar action of $\Lambda^*(H_1(Y; \mathbb{Z})/\operatorname{Tors})$ on singly pointed Heegaard Floer homology described in, e.g. [OS04b], though there are some key differences. For one thing, the maps A_λ in this section don't always square to zero, and they also don't always commute, even on homology. Nonetheless they are essential for the construction of graph cobordism maps. The work here is inspired by work in [Lip06] and [Ni10], and is technically very similar to the latter, where a very similar homological action is constructed on sutured Floer homology.

5.1. **Construction of the relative homology actions.** Fix a diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ for $(Y, \mathbf{w})$. Suppose that $(K, \partial K) \subseteq (\Sigma \times I \times \mathbb{R}, \mathbf{w} \times \{0, 1\} \times \mathbb{R})$ is an embedded surface which is $\mathbb{R}$-invariant. The case where $\partial K = \varnothing$ is allowed. Suppose further that the image $\pi_\Sigma(K)$ doesn't contain any intersection points between $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves and that $K \cap (\boldsymbol{\alpha} \cup \boldsymbol{\beta}) \times \partial I \times \mathbb{R}$ is trivial. One could try to define the map

$$A_K : CF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s}) \to CF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s})$$

by

$$A_K(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} (\#K\cap \operatorname{im}\phi)\cdot \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot \mathbf{y},$$

where $\#K\cap\operatorname{im}\phi$ is defined as the algebraic intersection number of K and $\operatorname{im}\phi$ as subsurfaces in $W_{\mathbb{D}} = \Sigma\times\mathbb{D}$ with boundaries fixed in

$$((\boldsymbol{\alpha}\cup\boldsymbol{\beta})\times(e_\alpha\cup e_\beta))\cup(\{\mathbf{w}\}\times\{0,1\}\times\mathbb{R}).$$

In the following motivating example, let $\rho : I\to[0,1/2)$ be a smooth function which $\rho(0)=\rho(1)=0$ and $\rho(t)>0$ for $t\notin\{0,1\}$.

Example 5.1. *If $\lambda : I\to\Sigma$ is a smooth immersed path from one basepoint w to another basepoint w', we define the following surface:*

$$K_\lambda^{\alpha\alpha} = \{(\lambda(t), 1-\rho(t), s) : t\in I, s\in\mathbb{R}\}$$

The map $A_{K_\lambda^{\alpha\alpha}}$ is one way to define the relative homology action.

The map $A_{K_\lambda^{\alpha\alpha}}$ defines a map for the path λ, but in proving invariance and relations involving $A_{K_\lambda^{\alpha\alpha}}$ we run into tranvsersality issues involving holomorphic curves with marked points. Instead we will define a map A_λ which should agree with $A_{K_\lambda^{\alpha\alpha}}$, but is defined more combinatorially. Our approach to defining the relative homology maps is essentially the same as the one taken in [Ni10] in defining an action of relative homology on sutured Floer homology. Note that since $A_{K_\lambda^{\alpha\alpha}}$ counts the algebraic intersection number, the choice of ρ does not affect the value of A_K. As a result, by letting ρ approach the function which is constantly zero, we obtain the following alternate characterization.

Suppose that λ intersects the curves $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ generically, i.e. it intersects each transversely and doesn't intersect any point in any $\alpha_i\cap\beta_j$. Let $p_1,\dots,p_m$ denote the points where λ intersects an $\boldsymbol{\alpha}$ curve. Given a homology disk ϕ, let $d_i(\lambda,\phi)$ denote the difference between the multiplicities of ϕ on the two sides of p_i. Define

$$a(\lambda,\phi) = \sum_{i=1}^m d_i(\lambda,\phi).$$

Then define

$$A_\lambda(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} a(\lambda,\phi)\cdot \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot \mathbf{y}.$$

Example 5.2. *If $\gamma : I\to\Sigma$ is a closed curve, then let $\widehat{\gamma} : I\to\Sigma\times(0,1)$ be an embedded closed curve which is homotopic to $\gamma\times\{1/2\} : I\to\Sigma\times I$. We define the surface*

$$K_\gamma = \{(\widehat{\gamma}(t), s), t\in I, s\in\mathbb{R}\}.$$

The map A_{K_γ} was used in the definition of the $H_1(Y;\mathbb{Z})$ action on $HF^\circ(Y,w,\mathfrak{s})$ for singly pointed diagrams (cf. [Lip06, Prop. 8.6.]*).*

As we did with the reltive homology maps, we can define maps A_γ more combinatorially by first defining $a(\gamma,\phi)$ as before, and then defining

$$A_\gamma(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} a(\gamma,\phi)\cdot \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot \mathbf{y}.$$

In principle $A_\gamma = A_{K_\gamma}$ and $A_\lambda = A_{K_\lambda}$, though it is easier to work with the maps A_γ than A_{K_γ}, so we will use A_γ and A_λ as the definitions of the (relative) homology maps.

Maps of the form A_K are often not chain maps. For the relative homology maps defined above we have the following:

Lemma 5.3. *Suppose that $\lambda : I\to\Sigma$ is an immersed path from w to w' which intersects the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves transversely and has trivial intersection with any $\alpha_i\cap\beta_j$. Then we have the following relation:*

$$A_\lambda\partial + \partial A_\lambda = U_w + U_{w'}.$$

If γ is a closed curve in Σ, then

$$A_\gamma \partial + \partial A_\gamma = 0.$$

Proof. We will show the first claim about A_λ and leave the necessary modifications for A_γ to the reader. This is analogous to [OS04b, Lem. 4.18]. Suppose that $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ is a homology disk with $\mu(\phi) = 2$. We will sum over the ends of $\widehat{\mathcal{M}}(\phi)$. There are an even number of ends, and hence

$$a(\lambda, \phi) \#(\partial \widehat{\mathcal{M}}(\phi)) = 0.$$

If $\mathbf{x} \neq \mathbf{z}$ then boundary degenerations cannot occur, so the ends of $\widehat{\mathcal{M}}(\phi)$ have the following form:

$$\partial(\widehat{\mathcal{M}}(\phi)) = \coprod_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}), \phi_2 \in \pi_2(\mathbf{y}, \mathbf{z}) \\ \mu(\phi_1) = \mu(\phi_2) = 1 \\ \phi_1 * \phi_2 = \phi}} \widehat{\mathcal{M}}(\phi_1) \times \widehat{\mathcal{M}}(\phi_2).$$

Note that if ϕ_1 and ϕ_2 are homology disks, then we have

$$a(\lambda, \phi_1 + \phi_2) = a(\lambda, \phi_1) + a(\lambda, \phi_2).$$

Combining these observations we have

$$0 = (a(\lambda, \phi_1) + a(\lambda, \phi_2)) \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}), \phi_2 \in \pi_2(\mathbf{y}, \mathbf{z}) \\ \mu(\phi_1) = \mu(\phi_2) = 1 \\ \phi_1 * \phi_2 = \phi}} \#\widehat{\mathcal{M}}(\phi_1) \#\widehat{\mathcal{M}}(\phi_2) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1) + n_{\mathbf{w}}(\phi_2)} \cdot \mathbf{z}.$$

Summing over all $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ with $\mu(\phi) = 2$, we get that the $\mathbf{z}$ component of $(A_\lambda \partial + \partial A_\lambda)(\mathbf{x})$ is zero.

In the case that $\mathbf{x} = \mathbf{z}$, we have extra contributions from $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ which degenerate into boundary degenerations. We have the additional counts of boundary degenerations, which add in a count of

$$\sum_{\substack{\phi \in \pi_2^\alpha(\mathbf{x}) \\ \mu(\phi) = 2}} a(\lambda, \phi) \#\widehat{\mathcal{N}}^\alpha(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{x} + \sum_{\substack{\phi \in \pi_2^\beta(\mathbf{x}) \\ \mu(\phi) = 2}} a(\lambda, \phi) \#\widehat{\mathcal{N}}^\beta(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{x}.$$

If $\phi \in \pi_2^\alpha(\mathbf{x})$ then ϕ can be represented as a sum of domains A_i where A_i is a connected component of $\Sigma \setminus \boldsymbol{\alpha}$. Similarly if $\phi \in \pi_2^\beta(\mathbf{x})$ then the domain of ϕ can be represented as a sum of the domains B_i, where B_i is a connected component of $\Sigma \setminus \boldsymbol{\beta}$. In light of the Maslov formula in Equation 3, we know that a nonnegative ϕ in $\pi_2^\alpha(\mathbf{x})$ or $\pi_2^\beta(\mathbf{x})$ with Maslov index 2 is of the form A_i or B_i for some i. On the other hand, it is clear from the definition that $a(\lambda, \phi) = 0$ unless $\phi = A_w$ or $\phi = A_{w'}$, where A_w and $A_{w'}$ are the components of $\Sigma \setminus \boldsymbol{\alpha}$ containing w or w' respectively. We can apply [OS08, Thm. 5.5], which shows that since we have multiple basepoints, we have $\#\widehat{\mathcal{N}}^\alpha(\phi) = 1 \in \mathbb{Z}_2$ if ϕ is a Maslov index 2 element of $\pi_2^\alpha(\mathbf{x})$.

In particular, summing over all $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ with Maslov index 2, we have that the $\mathbf{x}$ component of

$$(A_\lambda \partial + \partial A_\lambda + U_w + U_{w'})(\mathbf{x})$$

is zero. Combining this with the computation in the case that $\mathbf{x} \neq \mathbf{z}$, we get the desired formula. □

5.2. **Naturality of the relative homology maps.** In this section we prove that the relative homology maps A_λ are independent of the choice of diagram and the representative of the homotopy class of curve λ.

Lemma 5.4. *Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is strongly $\mathfrak{s}$-admissible. Then*

$$F^\circ_{\alpha\beta\gamma, \mathfrak{s}} \circ (A_\lambda \otimes 1) \simeq A_\lambda \circ F^\circ_{\alpha\beta\gamma, \mathfrak{s}}.$$

Proof. For a homology triangle ψ, define $a(\lambda, \psi) = \sum d_i(\psi)$ where $d_i(\psi)$ is the difference between the multiplicities of $\mathcal{D}(\psi)$ at each p_i, where $p_1, \ldots, p_m$ are the intersection points of the curve λ with $\boldsymbol{\alpha}$ curves.

It is clear that if ψ is a triangle and ϕ is a disk, then $a(\lambda, \psi + \phi) = a(\lambda, \psi) + a(\lambda, \phi)$. Now given a homology class $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ with $\mu(\psi) = 1$, we sum over the ends of $\mathcal{M}(\psi)$, which has a compactification as a compact 1-manifold. Since any nonconstant boundary degeneration or closed curve has Maslov index

at least two, we know that the only degenerations possible correspond to strip breaking, so the ends are modeled on

$$\coprod_{\substack{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta\\ \phi\in\pi_2(\mathbf{x},\mathbf{w}),\psi'\in\pi_2(\mathbf{w},\mathbf{y},\mathbf{z})\\ \mu(\phi)=1,\mu(\psi')=0\\ \phi+\psi'=\psi}} \widehat{\mathcal{M}}(\phi)\times\mathcal{M}(\psi') \sqcup \coprod_{\substack{\mathbf{w}\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma\\ \phi\in\pi_2(\mathbf{y},\mathbf{w}),\psi'\in\pi_2(\mathbf{x},\mathbf{w},\mathbf{z})\\ \mu(\phi)=1,\mu(\psi')=0\\ \phi+\psi'=\psi}} \widehat{\mathcal{M}}(\phi)\times\mathcal{M}(\psi') \tag{4}$$

$$\sqcup \coprod_{\substack{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma\\ \phi\in\pi_2(\mathbf{w},\mathbf{z}),\psi'\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{z})\\ \mu(\phi)=1,\mu(\psi')=0\\ \phi+\psi'=\psi}} \mathcal{M}(\psi')\times\widehat{\mathcal{M}}(\phi).$$

Define the map $H^\lambda_{\alpha\beta\gamma,\mathfrak{s}}$ by

$$H^\lambda_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) = \sum_{\substack{\mathbf{z}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma\\ \psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{z})\\ \mu(\psi)=0}} a(\lambda,\psi)\#\mathcal{M}(\psi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)}\cdot\mathbf{z}.$$

Since the total count of ends in Equation (4) is zero, multiplying the total count by $a(\lambda,\psi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)}$, and then summing over all $\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{z})$ with Maslov index one, we get that the $\mathbf{z}$ coefficient of

$$F^\circ_{\alpha\beta\gamma,\mathfrak{s}}(A_\lambda(\mathbf{x})\otimes\mathbf{y}) + A_\lambda(F^\circ_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y})) + \partial_{\alpha\gamma}H^\lambda_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) + H^\lambda_{\alpha\beta\gamma,\mathfrak{s}}((\partial_{\alpha\beta}\mathbf{x})\otimes\mathbf{y}+\mathbf{x}\otimes(\partial_{\beta\gamma}\mathbf{y}))$$

is equal to zero. Hence the desired relation follows. □

Lemma 5.5. *If J_s and J'_s are two almost complex structure on $W_{\mathbb{D}}$ which satisfy axioms $(J1)-(J6)$ in* [Lip06]*, then*

$$A_\lambda\Phi_{J_sJ'_s}\simeq\Phi_{J_sJ'_s}A_\lambda,$$

where $\Phi_{J_sJ'_s}$ is the change of almost complex structures map.

Proof. Pick an almost complex structure $\widehat{J}$ on $W_{\mathbb{D}}$ which agrees with J_s on $\Sigma\times I\times[1,\infty)$ and with J'_s on $\Sigma\times I\times(-\infty,-1]$, and which satisfies the required axioms in [Lip06] to compute the map $\Phi_{J_sJ'_s}$. The map $\Phi_{J_sJ'_s}$ is computed by counting $\widehat{J}$-holomorphic disks of Maslov index zero.

We define a map $H_{\widehat{J},\lambda}$ by the formula

$$H_{\widehat{J},\lambda}(\mathbf{x}) = \sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta\\ \phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=0}} a(\lambda,\phi)\#\mathcal{M}_{\widehat{J}}(\phi).$$

Counting the ends of a moduli space of the form $\mathcal{M}_{\widehat{J}}(\phi)$ where $\mu(\phi)=1$, and then summing over all $\phi\in\pi_2(\mathbf{x},\mathbf{y})$ yields that the $\mathbf{y}$ coefficient of

$$(A_\lambda\Phi_{J_sJ'_s}+\Phi_{J_sJ'_s}A_\lambda+\partial H_{\widehat{J},\lambda}+H_{\widehat{J},\lambda}\partial)(\mathbf{x})$$

is zero. □

Lemma 5.6. *Suppose that λ and λ' are two paths in Σ starting and ending at the basepoints w and w' which are homotopic relative their endpoints. Then*

$$A_\lambda\simeq A_{\lambda'}.$$

Proof. The homotopy between λ and λ' can be replaced by a sequence of homotopies of the following form

(1) homotopies which don't affect the intersection of λ which any $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves;
(2) homotopies which push a piece of λ through a point on $\boldsymbol{\beta}\setminus\boldsymbol{\alpha}$;
(3) homotopies which push a piece of λ through a point on $\boldsymbol{\alpha}\setminus\boldsymbol{\beta}$;
(4) homotopies which push a piece of λ across an intersection point of $\boldsymbol{\alpha}\cap\boldsymbol{\beta}$.

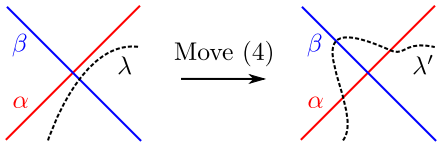

FIGURE 3. Move (4) in Lemma 5.6

Homotopies of the first two types trivially don't affect $a(\lambda, \phi)$. It is an easy computation to see that move (3) also doesn't affect $a(\lambda, \phi)$. Move (4) is shown in Figure 3.

Let p be the intersection point across which we are pushing λ. Let a, b, c, d be the multiplicities, labeled clockwise, as in Figure 4.

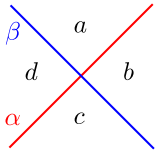

FIGURE 4. Multiplicities around the vertex in Move(4).

We compute that

$$a(\lambda', \phi) - a(\lambda, \phi) = (b - a) + (d - c).$$

Combining this with the vertex relations in Equation (1), we see

$$a(\lambda', \phi) - a(\lambda, \phi) = \delta(\mathbf{y}, p) - \delta(\mathbf{x}, p).$$

Thus we have that

$$\begin{aligned} A_{\lambda'}(\mathbf{x}) - A_{\lambda}(\mathbf{x}) &= \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 0}} (\delta(\mathbf{y}, p) - \delta(\mathbf{x}, p)) \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y} \\ &= H_p(\partial \mathbf{x}) - \partial(H_p(\mathbf{x})), \end{aligned}$$

where $H_p(\mathbf{z}) = \delta(\mathbf{z}, p) \cdot \mathbf{z}$, yielding the desired chain homotopy. □

Lemma 5.7. *Suppose that λ is a path from w_1 to w_2 and λ' is a path from w_2 to w_3. Then*

$$A_{\lambda' * \lambda} = A_\lambda + A_{\lambda'}.$$

Proof. This follows directly from the formula since $a(\lambda' * \lambda, \phi) = a(\lambda', \phi) + a(\lambda, \phi)$. □

Similarly, we also have the following:

Lemma 5.8. *If λ is a path from w to w', and γ and γ' are two closed loops which start and end at w, then we have*

$$A_{\gamma' * \gamma} = A_{\gamma'} + A_\gamma$$

and also

$$A_\gamma + A_\lambda = A_{\gamma * \lambda}.$$

Proof. This follows as before since $a(\lambda' * \lambda, \phi) = a(\lambda', \phi) + a(\lambda, \phi)$. □

In contrast to the hat case, the relative homology maps do not commute. Instead we have the following:

Lemma 5.9. *Suppose λ_1 and λ_2 are two paths between basepoints. If $E(\lambda_i)$ denotes the endpoints of λ_i, then we have the following:*

$$A_{\lambda_1} A_{\lambda_2} + A_{\lambda_2} A_{\lambda_1} = \sum_{w \in E(\lambda_1) \cap E(\lambda_2)} U_w.$$

Proof. We proceed similarly to Lemma 5.3. We wish to count the ends of $\widehat{\mathcal{M}}(\phi)$ where $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ is a Maslov index 2 disk. There are two cases to consider: namely $\mathbf{x} \neq \mathbf{z}$ and $\mathbf{x} = \mathbf{z}$.

In the case that $\mathbf{x} \neq \mathbf{z}$, the ends of $\widehat{\mathcal{M}}(\phi)$ for a Maslov index two ϕ correspond only to strip breaking into two Maslov index one strips. Summing over all such ϕ, we have

$$0 = \sum_{\substack{\phi \in \pi_2(\mathbf{x},\mathbf{z}) \\ \mu(\phi)=2}} a(\lambda_1, \phi)a(\lambda_2, \phi) \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x},\mathbf{y}), \phi_2 \in \pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1 \\ \phi_1+\phi_2=\phi}} \#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{z}.$$

Computing that

$$a(\lambda_1, \phi_1 + \phi_2)a(\lambda_2, \phi_1 + \phi_2)$$

$$= \Big(a(\lambda_1, \phi_1)a(\lambda_2, \phi_2) + a(\lambda_1, \phi_2)a(\lambda_2, \phi_1)\Big) + \Big(a(\lambda_1, \phi_1)a(\lambda_2, \phi_1) + a(\lambda_1, \phi_2)a(\lambda_2, \phi_2)\Big),$$

we can rewrite the above expression to yield

$$0 = \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x},\mathbf{y}), \phi_2 \in \pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda_1, \phi_1)a(\lambda_2, \phi_2) + a(\lambda_1, \phi_2)a(\lambda_2, \phi_1))\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1)+n_{\mathbf{w}}(\phi_2)} \cdot \mathbf{z}$$

$$+ \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x},\mathbf{y}), \phi_2 \in \pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda_1, \phi_1)a(\lambda_2, \phi_1) + a(\lambda_1, \phi_2)a(\lambda_2, \phi_2))\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1)+n_{\mathbf{w}}(\phi_2)} \cdot \mathbf{z}.$$

The above expression is exactly the $\mathbf{z}$ component of

$$(A_{\lambda_1}A_{\lambda_2} + A_{\lambda_2}A_{\lambda_1} + \partial H_{\lambda_1\lambda_2} + H_{\lambda_1\lambda_2}\partial)(\mathbf{x}),$$

where $H_{\lambda_1\lambda_2}$ is the map defined by

$$H_{\lambda_1\lambda_2}(\mathbf{x}) = \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi \in \pi_2(\mathbf{x},\mathbf{y}) \\ \mu(\phi)=1}} a(\lambda_1, \phi)a(\lambda_2, \phi)\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}.$$

In the case that $\mathbf{x} = \mathbf{z}$, we have the additional possibility of boundary degenerations. These make an additional contribution of

$$\sum_{\substack{\phi \in \pi_2^\alpha(\mathbf{x}) \\ \mu(\phi)=2}} a(\lambda_1, \phi)a(\lambda_2, \phi)\#\widehat{\mathcal{N}}^\alpha(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{x} + \sum_{\substack{\phi \in \pi_2^\beta(\mathbf{x}) \\ \mu(\phi)=2}} a(\lambda_1, \phi)a(\lambda_2, \phi)\#\widehat{\mathcal{N}}^\beta(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{x}.$$

We now note that $a(\lambda, \phi) = 0$ whenever $\phi \in \pi_2^\beta(\mathbf{x})$ and also that $a(\lambda_1, \phi)a(\lambda_2, \phi)$ is zero unless $\mathcal{D}(\phi) = A_w$ for some $w \in E(\lambda_1) \cap E(\lambda_2)$, in which case it is 1. The desired formula thus follows. □

Lemma 5.10. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for Y and that γ is an immersed, closed curve in Σ which intersects no points of $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$ and also intersects the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves transversely. If γ is homotopic to an $\boldsymbol{\alpha}$ curve, or to a $\boldsymbol{\beta}$ curve, then $A_\gamma \simeq 0$.*

Proof. By Lemma 5.6 a homotopy of γ results in a chain homotopic A_γ. If γ is homotopic to an $\boldsymbol{\alpha}$ curve, then it is also homotopic to a small push off of a $\boldsymbol{\alpha}$ curve since Σ is orientable. Such a small push off does not intersect any $\boldsymbol{\alpha}$ curves, and hence the quantity $a(\gamma, \phi) = 0$ for all ϕ since $a(\gamma, \phi)$ is an empty sum. Hence $A_\gamma = 0$.

If γ is homotopic to a $\boldsymbol{\beta}$ curve, then by pushing off of the $\boldsymbol{\beta}$ curve, we can assume that γ doesn't intersect any $\boldsymbol{\beta}$ curve. In that case, $a(\gamma, \phi) = 0$ since $a(\gamma, \phi)$ is a sum of differences which telescope since γ doesn't cross any $\boldsymbol{\beta}$ curves. In particular $A_\gamma = 0$. □

In the context of a single basepoint in Heegaard Floer homology, the homology action squares to zero (cf. [OS04b, Prop. 4.17], [Lip06, Prop. 8.6]). If we were to proceed to prove the following lemma using those techniques, we would need to prove a transversality result for curves with two marked points, though it's not clear if such a transversality result even holds in our context. Our proof of the following lemma requires no such transversality result:

Lemma 5.11. *Suppose that λ is an immersed path from w to w' in Σ and γ is an immersed, closed loop in Σ. Then*

$$A_\lambda \circ A_\lambda \simeq U_w \simeq U_{w'}$$

and

$$A_\gamma \circ A_\gamma \simeq 0.$$

Proof. Note that $U_w \simeq U_{w'}$ by Lemma 5.3. We will prove that $A_\lambda \circ A_\lambda \simeq U_w$. In constructing the chain homotopy, there is a place where we need to pick a sign. One choice of the sign yields a homotopy to U_w and the other choice yields a chain homotopy to $U_{w'}$. The proof of $A_\lambda \circ A_\lambda \simeq 0$ follows with only minor modification. First pick an orientation of λ. As usual we will count the ends of moduli spaces. Let $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ be a homotopy class of disks of Maslov index 2 and first suppose that $\mathbf{x} \neq \mathbf{z}$. Pick a coherent orientation system as in [Lip06, Sec. 6.]. This gives an orientation the space $\widehat{\mathcal{M}}(\phi)$, a 1-manifold with boundary. Since $\mathbf{x} \neq \mathbf{z}$, the boundary corresponds to exactly strip breaking. We have that

$$\partial(\widehat{\mathcal{M}}(\phi)) = \coprod_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1 \\ \phi_1*\phi_2=\phi}} \widehat{\mathcal{M}}(\phi_1)\times\widehat{\mathcal{M}}(\phi_2)$$

and hence

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2) = 0.$$

Here the above expression is an equation in $\mathbb{Z}$ since we picked an orientation on $\widehat{\mathcal{M}}(\phi)$. Since we also picked an orientation of λ, we can take the expression $a(\lambda,\phi)$ to have values in $\mathbb{Z}$. Multiplying the above expression by $a(\lambda,\phi)^2$, we get

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda,\phi_1)+a(\lambda,\phi_2))^2\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2) = 0,$$

which yields that

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda,\phi_1)^2+a(\lambda,\phi_2)^2)\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2) + 2 \sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda,\phi_1)a(\lambda,\phi_2))\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2) = 0.$$

Now add

$$a(\lambda,\phi) \sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2) = 0$$

to the left side of the above equation, then divide the entire expression by 2 to see that

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \left(\frac{a(\lambda,\phi_1)(a(\lambda,\phi_1)+1)}{2} + \frac{a(\lambda,\phi_2)(a(\lambda,\phi_2)+1)}{2}\right)\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)$$

$$+ \sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda,\phi_1)a(\lambda,\phi_2))\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2) = 0.$$

Multiply both sides by $U_{\mathbf{w}}^{n_{\mathbf{w}}}(\phi)$ and then sum over all $\phi \in \pi_2(\mathbf{x},\mathbf{z})$ with Maslov index 2, then project all coefficients to $\mathbb{Z}_2$ and we get that the $\mathbf{z}$ component of

$$(\partial H_\lambda + H_\lambda \partial + A_\lambda \circ A_\lambda)(\mathbf{x}),$$

is zero, where

$$H_\lambda(\mathbf{x}) = \sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi\in\pi_2(\mathbf{x},\mathbf{y}) \\ \mu(\phi)=1}} \left(\frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}\right)\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{y}.$$

Here of course the expression $(a(\lambda,\phi_2)(a(\lambda,\phi_2)+1)/2$ depends on the orientation of λ and is first computed in $\mathbb{Z}$, then projected to $\mathbb{Z}_2$.

In the case that $\mathbf{x}=\mathbf{z}$, as before, $\#\widehat{\mathcal{M}}(\phi)$ now has additional boundary components corresponding to boundary degenerations. The strategy is similar. Proceeding as before, we get

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda,\phi_1)^2+a(\lambda,\phi_2)^2)\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)+2\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda,\phi_1)a(\lambda,\phi_2))\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)$$

$$+a(\phi,\lambda)^2\#\widehat{\mathcal{N}}^\alpha(\phi)+a(\phi,\lambda)^2\#\widehat{\mathcal{N}}^\beta(\phi)=0.$$

Now add

$$a(\lambda,\phi)\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)+a(\lambda,\phi)\#\widehat{\mathcal{N}}^\alpha(\phi)+a(\lambda,\phi)\#\widehat{\mathcal{N}}^\beta(\phi)=0,$$

to the left side, and divide by 2 to get

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \left(\frac{a(\lambda,\phi_1)(a(\lambda,\phi_1)+1)}{2}+\frac{a(\lambda,\phi_2)(a(\lambda,\phi_2)+1)}{2}\right)\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)$$

$$+\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} (a(\lambda,\phi_1)a(\lambda,\phi_2))\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)$$

$$+\frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}\#\widehat{\mathcal{N}}^\alpha(\phi)+\frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}\#\widehat{\mathcal{N}}^\beta(\phi)=0.$$

Now multiply by $U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}$ and sum over all $\phi\in\pi_2(\mathbf{x},\mathbf{x})$. This yields that the $\mathbf{x}$ component of

$$(\partial H_\lambda + H_\lambda \partial + A_\lambda \circ A_\lambda)(\mathbf{x})$$

$$+\sum_{\substack{\phi\in\pi_2^\alpha(\mathbf{x}) \\ \mu(\phi)=2}} \frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}\#\widehat{\mathcal{N}}^\alpha(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x}+\sum_{\substack{\phi\in\pi_2^\beta(\mathbf{x}) \\ \mu(\phi)=2}} \frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}\#\widehat{\mathcal{N}}^\beta(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x},$$

is zero. Note that if A_w and $A_{w'}$ denote the Maslov index 2 homology class in $\pi_2^\alpha(\mathbf{x})$ determined by unique component of $\Sigma\setminus\boldsymbol{\alpha}$ containing w (resp. w'), then

$$\sum_{\substack{\phi\in\pi_2^\alpha(\mathbf{x}) \\ \mu(\phi)=2}} \frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}(\#\widehat{\mathcal{N}}^\alpha(\phi))U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x}+\sum_{\substack{\phi\in\pi_2^\beta(\mathbf{x}) \\ \mu(\phi)=2}} \frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}(\#\widehat{\mathcal{N}}^\beta(\phi))U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x}$$

$$=\pm\frac{a(\lambda,A_w)(a(\lambda,A_w)+1)}{2}U_w\cdot\mathbf{x}\pm\frac{a(\lambda,A_{w'})(a(\lambda,A_{w'})+1)}{2}U_{w'}\cdot\mathbf{x}.$$

Note that exactly one of $a(\lambda,A_w)$ and $a(\lambda,A_{w'})$ is $+1$, and the other is -1, without loss of generality, say that $a(\lambda,A_w)=1$. Hence

$$\pm\frac{a(\lambda,A_w)(a(\lambda,A_w)+1)}{2}U_w\cdot\mathbf{x}\pm\frac{a(\lambda,A_{w'})(a(\lambda,A_{w'})+1)}{2}U_{w'}\cdot\mathbf{x}=\pm U_w\cdot\mathbf{x}.$$

Hence we conclude that

$$(\partial H_\lambda + H_\lambda \partial + A_\lambda \circ A_\lambda + U_w)(\mathbf{x}) = 0,$$

from which the claim follows.

The proof easily adapts to compute $A_\gamma \circ A_\gamma \simeq 0$. In that case, $a(\gamma, A) = 0$ for any homotopy class of boundary degeneration A, so boundary degenerations have no contribution. □

We state the following fact, though we will not need it in this paper:

Lemma 5.12. *If γ is a closed loop in Σ representing an element of* Tors $\subseteq H_1(Y)$*, then $A_\gamma \simeq 0$.*

Proof. This can be proven using the same argument in [Ni10, Lem. 2.4], using only a change of notation. □

6. Free stabilization maps

In this section we introduce maps associated to adding or removing a basepoint. These correspond to the *creation/termination* maps defined by the author in [Zem15] for the hat flavor. Suppose that $w \notin \mathbf{w}$ is a basepoint, and $(\sigma', \mathfrak{P})$ is a coloring on $\mathbf{w} \cup \{w\}$, which extends the coloring $(\sigma, \mathfrak{P})$ on $\mathbf{w}$. We will define homomorphisms of $\mathbb{Z}_2[U_{\mathfrak{P}}]$-modules

$$S_w^+ : HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S_w^- : HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

which we will call the **free stabilization maps**.

We now summarize these maps. Given an embedded diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ with is strongly $\mathfrak{s}$-admissible, we say that $\mathcal{H}$ is **subordinate to** the point w if $w \in \Sigma$ and w does not intersect any of the $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves. Now given a collection of data $\mathcal{J}$ (which we will describe more precisely below), we can form a diagram $\mathcal{H}_\mathcal{J}$

$$\mathcal{H}_\mathcal{J} = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w} \cup \{w\})$$

where α_0 and β_0 are two small isotopic curves each bounding a small ball centered around the point w, which don't intersect any of the other $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves and such that $|\alpha_0 \cap \beta_0| = 2$. The two intersection points of $\alpha_0 \cap \beta_0$ are distinguished by the Maslov index. We let θ^+ and θ^- denote the higher and lower graded intersection points respectively.

For appropriate almost complex structures, we will define the maps $S_w^\pm$ by the formulas

$$S_w^+(\mathbf{x}) = \mathbf{x} \times \theta^+,$$

and

$$S_w^-(\mathbf{x} \times \theta^-) = \mathbf{x}, \qquad \text{and} \qquad S^-(\mathbf{x} \times \theta^+) = 0.$$

The rest of the section is devoted to proving Theorem 6.22, which shows that the free stabilization maps are natural.

6.1. Gluing data, neck stretching, and gluing moduli spaces.

To define the free stabilization maps, as well as the 1- and 3-handle maps, we will need to perform neck stretching and gluing arguments repeatedly. We summarize the procedure in this subsection, and in the following subsection we will embark upon applying it to analyze the free stabilization maps. Consider the diagram $(S^2, \alpha_0, \beta_0, w, w')$ for (S^3, w, w') shown in Figure 5.

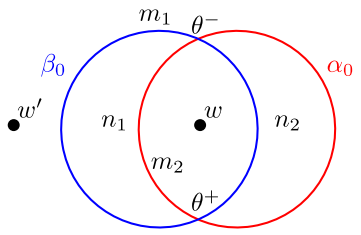

FIGURE 5. The standard diagram for (S^3, w, w'), with multiplicities labelled. The point p' is chosen in the same domain as w'.

We make the following definition:

Definition 6.1. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(Y, \mathbf{w})$ which is subordinate to a point $w \in Y \setminus \mathbf{w}$. We call a collection $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ **gluing data**, if*

1. *B is a closed ball containing w on Σ which doesn't intersect any $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves;*
2. *B' is a closed ball on S^2 which doesn't intersect α_0 or β_0 in the diagram $(S^2, \alpha_0, \beta_0, w, w')$ shown in Figure 5 such that B' contains the point $p' \in S^2$;*
3. *J_s is an almost complex structure on $\Sigma \times \mathbb{D}$ which is split on B;*
4. *J_s' is an almost complex structure on $S^2 \times \mathbb{D}$ which is split on B';*
5. *r and r' are real numbers such that $0 < r, r' < 1$;*
6. *using the unique (up to rotation) conformal identifications of (B, w) and (B', p') as $(\mathbb{D}, 0)$, where $\mathbb{D}$ denotes the unit complex disk, ι is an embedding of $\overline{S^2 \setminus r' \cdot B'}$ into $r \cdot B \subseteq \Sigma$ such that*
$$\overline{r \cdot B} \setminus \iota(S^2 \setminus r' \cdot B')$$
is a closed annulus.
7. *letting $\widetilde{A}, A$ and A' denote the closed annuli $B \setminus \iota(S^2 \setminus B')$, $B \setminus r \cdot B$ and $B' \setminus r' \cdot B'$ respectively, then*
$$\phi : \widetilde{A} \to S^1 \times [-a, 1+b]$$
is a diffeomorphism which sends the annulus A to $[-a, 0]$ and $\iota(A')$ to $[1, 1+b]$ and is conformal on A and A'.

Given a diagram $\mathcal{H}$ with gluing data $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ and a $T > 1$, we can form the diagram

$$\mathcal{H}_{\mathcal{J}} = (\Sigma, \boldsymbol{\alpha} \cup \{\iota\alpha_0\}, \boldsymbol{\beta} \cup \{\iota\beta_0\}, \mathbf{w} \cup \{w_0\}),$$

and form an almost complex structure $\mathcal{J}(T)$ on $\mathcal{H}_{\mathcal{J}}$. The almost complex structure $\mathcal{J}(T)$ is defined to be J_s on $\Sigma \setminus (r \cdot B)$, and $\iota_* J_s'$ on $\iota(S^2 \setminus (r' \cdot B'))$. The map ϕ can be modified to a map $\phi_T : \widetilde{A} \to S^1 \times [-a, T+b]$ which is conformal on A and A' by stretching in the region $S^1 \times [0, 1]$. This can be done in an isotopically unique way, and the dependence is explored in Lemma 6.7. The complex structure $\mathcal{J}(T)$ is defined to be $(\phi_T)^* J_{S^1 \times [-a, T+b]}$.

The free stabilization maps will then be maps

$$S_w^+ : CF^\circ_{J_s}(\mathcal{H}) \to CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}})$$

and

$$S_w^- : CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}}) \to CF^\circ_{J_s}(\mathcal{H}),$$

for sufficiently large T. The phrase "sufficiently large" is given a precise meaning in Definition 6.9.

As we let T approach $+\infty$, the almost complex structure $\mathcal{J}(T)$ approaches $J_s \vee J_s'$ on $(\Sigma \vee_{w=p'} S^2) \times I \times \mathbb{R}$. If u_{T_n} is a sequence of $\mathcal{J}(T_n)$-holomorphic disks on $\mathcal{H}_{\mathcal{J}}$, representing a homology class $\phi = \phi_1 \# \phi_2$, we can extract a Gromov limit of u_{T_n} to a broken holomorphic disk U_1 on $\Sigma \times I \times \mathbb{R}$ representing ϕ_1 and a broken holomorphic disk U_2 on $\Sigma_2 \times I \times \mathbb{R}$ representing ϕ_2. By a broken holomorphic disk U_i, we mean a collection of holomorphic curves (genuine holomorphic disks, boundary degenerations, closed curves) whose total homology class represents ϕ_i. In general such a collection may be quite complicated, but in the case that both U_1 and U_2 are each genuine holomorphic disks, u_1 and u_2, there is a matching condition which must be satisfied:

$$\rho^{p_1}(u_1) = \rho^{p_2}(u_2),$$

where $\rho^{p_i}(u_i)$ is the divisor in $\mathrm{Sym}^{n_{p_i}(u_i)}(\mathbb{D})$ defined by

$$\rho^{p_i}(u_i) = (\pi_{\mathbb{D}} \circ u_i) \circ (\pi_{\Sigma_i} \circ u_i)^{-1}(p_i).$$

Conversely, if u_1 and u_2 are holomorphic disks on $\Sigma \times I \times \mathbb{R}$ and $S^2 \times I \times \mathbb{R}$ respectively which satisfy $\rho^{p_1}(u_1) = \rho^{p_2}(u_2)$, we say that u_1 and u_2 are **preglued**. In general, given a preglued pair u_1 and u_2, we can glue u_1 and u_2 to get a $\mathcal{J}(T)$-holomorphic curve $u_1 \# u_2$ for all sufficiently large T which achieve transversality (more precisely, under appropriate assumptions, there is an isomorphism of the moduli space $\mathcal{M}(\phi)$ with the fibered product $\mathcal{M}(\phi_1) \times_\rho \mathcal{M}(\phi_2)$). Technical details can be found in, for example, [Lip06, App. A], [OS08, Thm. 5.1], [OS04b], or more generally [BEH+03].

6.2. **Counting holomorphic strips after free stabilization.** In this subsection we begin our analysis of the free stabilization maps. We prove that the free stabilization maps are chain maps for sufficiently large T, and that they commute with the relative homology maps under appropriate circumstances.

Consider the diagram for $(S^2, \alpha_0, \beta_0, w, w')$ pictured in Figure 5. Let p' be a point in the region containing the basepoint w'. Given a homology disk ϕ and a divisor $\mathbf{d} \in \mathrm{Sym}^k(\mathbb{D})$, we define

$$\mathcal{M}(\phi, \mathbf{d}) = \{u \in \mathcal{M}(\phi) : \rho^{p_2}(u) = \mathbf{d}\}.$$

Ozsváth and Szabó prove the following using a Gromov compactness argument:

Lemma 6.2 ([OS08, Lem. 6.4])**.** *Let $(S^2, \alpha_0, \beta_0, w, w')$ be the doubly pointed diagram for (S^3, w, w') shown in Figure 5. If $\mathbf{d}$ is a generic divisor in $\mathrm{Sym}^k(\mathbb{D})$ for some integer $k > 0$, then*

$$\sum_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi)=2k \\ n_{w'}(\phi)=0}} \#\mathcal{M}(\phi, \mathbf{d}) = 1 \in \mathbb{Z}_2.$$

Lemma 6.3. *Suppose ϕ is a disk in the diagram $(S^2, \alpha_0, \beta_0, w, w')$ shown in Figure 5. If $\phi \in \pi_2(\theta^+, \theta^+)$ or $\phi \in \pi_2(\theta^-, \theta^-)$ then*

$$\mu(\phi) = 2m_1 + 2m_2.$$

If $\phi \in \pi_2(\theta^+, \theta^-)$ then

$$\mu(\phi) = 2m_1 + 2m_2 + 1.$$

If $\phi \in \pi_2(\theta^-, \theta^+)$, then

$$\mu(\phi) = 2m_1 + 2m_2 - 1.$$

Proof. Suppose that $\phi \in \pi_2(\theta^+, \theta^+)$ or $\phi \in \pi_2(\theta^-, \theta^-)$. Note that the stated formula holds if ϕ is a constant disk. Any other homotopy class is obtained by splicing in an α boundary degeneration or a β boundary degeneration, and the formula respects this.

Similarly for the case of $\phi \in \pi_2(\theta^+, \theta^-)$, we note that the formula holds for the disk with $n_1 = 1$ and all other multiplicities equal to zero. Any other disk can be obtained by splicing in boundary degenerations, and the formula respects that operation. The formula for $\phi \in \pi_2(\theta^-, \theta^+)$ follows similarly. □

Before we embark on the necessary gluing arguments needed to prove naturality, we first prove an admissibility result for stabilizing strongly $\mathfrak{s}$-admissible diagrams:

Lemma 6.4. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for $(Y, \mathbf{w})$ which is subordinate to $w \in Y \setminus \mathbf{w}$. Let $\mathcal{H}_\mathcal{J}$ be the diagram described above for $(Y, \mathbf{w} \cup \{w\})$ for some choice of gluing data $\mathcal{J}$. Suppose $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. Then $\mathcal{H}$ is strongly $\mathfrak{s}$-admissible iff $\mathcal{H}_\mathcal{J}$ is strongly $\mathfrak{s}$-admissible.*

Proof. Suppose $\mathcal{H}$ is strongly $\mathfrak{s}$-admissible. Suppose $\mathcal{D} = \mathcal{D}_1 \# \mathcal{D}_2$ is a periodic domain on $\mathcal{H}_\mathcal{J}$ with $\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 2N$, where $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ and $\mathcal{D}_2$ is a

$$\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = \langle c_1(\mathfrak{s}), H(\mathcal{D}_1)\rangle,$$

so $\mathcal{D}_1$ has a multiplicity greater than N, and hence so does $\mathcal{D}$. It follows that $\mathcal{H}_\mathcal{J}$ is strongly $\mathfrak{s}$-admissible.

Conversely if $\mathcal{H}_\mathcal{J}$ is strongly $\mathfrak{s}$-admissible and $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ with $\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 2N$. Let A_w and $A_{w'}$ be the components of $S^2 \setminus \alpha_0$ containing w and w' respectively. Let $\mathcal{D}_2$ be the domain on $(S^2, \alpha_0, \beta_0, w, w')$ defined by

$$\mathcal{D}_2 = n_w(\mathcal{D}_1) \cdot A_{w'}.$$

Since $\mathcal{H}_\mathcal{J}$ is strongly $\mathfrak{s}$-abmissible, we know that $\mathcal{D}_1 \# \mathcal{D}_2$ has coefficients strictly bigger than N, and hence so does $\mathcal{D}_1$, so $\mathcal{H}$ is strongly $\mathfrak{s}$-admissible. □

We now show that the free stabilization maps are chain maps:

Lemma 6.5. *Fix a Heegaard diagram $\mathcal{H}$ for $(Y, \mathbf{w})$ which is subordinate to $w \in Y \setminus \mathbf{w}$, as well as gluing data $\mathcal{J}$. If ∂_{J_s} denotes the differential of $CF^\circ_{J_s}(\mathcal{H}, \mathfrak{s})$ and $\partial_{\mathcal{J}(T)}$ denotes the differential of $CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_\mathcal{J}, \mathfrak{s})$, then for all sufficiently large T which obtain transversality, we have an identification:*

$$\partial_{\mathcal{J}(T)} = \begin{pmatrix} \partial_{J_s} & * \\ 0 & \partial_{J_s} \end{pmatrix}.$$

Proof. This is essentially [OS08, Prop. 6.5], though we will prove it here as well, since later arguments will be very analogous. Also note that in [OS08, Prop. 6.5], Ozsváth and Szabó show that if we let p' approach a point on $\beta_0 \setminus \alpha_0$, then we can take $* = U_w - U_{w'}$, though we will not have use for that at the moment.

Write

$$\partial_{\mathcal{J}(T)} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

We need to compute A, D and C. Let us first show that $A = D = \partial_{J_s}$ for sufficiently large T. Suppose that $\phi = \phi_1 \# \phi_2$ is a homology class of disks. Using the combinatorial formula for the Maslov index (cf. Eq. (2) or [Sar11a]), we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 = \mu(\phi_1) + 2m_2.$$

We first show that any disk with $\mathcal{J}(T_i)$-holomorphic representative for arbitrary large i corresponds to gluing a preglued flowline. Given a sequence of disks u_T representing $\phi = \phi_1 \# \phi_2$ for the almost complex structure $\mathcal{J}(T)$, we know that as $T \to \infty$, the holomorphic strips u_T have a Gromov limit to a broken disk U_1 on Σ and a broken disk U_2 on S^2. Since $\mu(\phi) = 1$, and $m_2 \geq 0$, using the Maslov index computation above, we conclude that $\mu(\phi_1) = 1$ and $m_2 = 0$, so U_1 consists of a single holomorphic strip u_1. There must be a component u_2 of U_2 which satisfies the matching condition $\rho^{p_1}(u_1) = \rho^{p_2}(u_2)$. Say that u_2 represents the homology class ϕ_2'. We wish to show that u_2 is the only nonconstant flowline in U_2. We know that $\phi_1 \# \phi_2'$ is nonconstant and has holomorphic representative, and hence we must have $\mu(\phi_1 \# \phi_2') \geq 1$. On the other hand, we know that

$$\mu(\phi_1 \# \phi_2) = \mu(\phi_1 \# \phi_2') + \mu(0 \# (\phi_2 - \phi_2')) = \mu(\phi_1 \# \phi_2') + \mu(\phi_2 - \phi_2'). \tag{5}$$

Since $\phi_2 - \phi_2'$ has holomorphic representatives as a broken disk (the remaining components of U_2), we conclude that

$$\mu(\phi_2 - \phi_2') \geq 0,$$

with equality iff $\phi_2 = \phi_2'$. On the other hand, in light of the fact that the quantity in Equation (5) is equal to 1, we conclude that $\phi_2 = \phi_2'$, and hence that u_2 represents all of the nonconstant components of ϕ_2.

In particular we conclude that any such sequence of u_T representing ϕ has a Gromov limit as a preglued flowline (u_1, u_2) with $m_2(u_2) = 0$ and $\mu(u_1) = 1$.

On the other hand, given a flowline u_1 on Σ representing ϕ_1, applying Lemma 6.2 we see that the total number of holomorphic strips u_2 representing a class $\phi_2 \in \pi_2(\theta^+, \theta^+)$ with Maslov index $2n_{p_1}(\phi_1)$ and $m_2(\phi_2) = 0$, with (u_1, u_2) preglued is $1 \in \mathbb{Z}_2$. The same statement holds for disks in $\pi_2(\theta^+, \theta^+)$.

Since we are working with strongly $\mathfrak{s}$-admissible diagrams, by Lemma 4.9, there are only finitely many homotopy classes of disks with nonnegative domain and Maslov index 1, so we can pick T large enough for the previous argument to apply to all of them simultaneously. Hence we conclude that $A = D = \partial_{J_s}$.

We finally need to show that $C = 0$. Given a disk $\phi = \phi_1 \# \phi_2$, we use the combinatorial formula for the Maslov index to see that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 = \mu(\phi_1) + 2m_2 + 1.$$

If ϕ has holomorphic representatives u_T for arbitrarily large T, then from a weak limit argument we conclude that $\mu(\phi_1) \geq 0$. Since $m_2 \geq 0$, we conclude that $\mu(\phi_1) = 0$. By stretching the neck sufficiently, we see that ϕ_1 has a representative as a broken holomorphic strip, and since $\mu(\phi_1) = 0$, we conclude that ϕ_1 must be the constant homology class. Hence we see that

$$n_w(\phi_1) = n_p(\phi_2) = 0.$$

Since $n_{p'}(\phi_2) = 0$ as well, and $\phi_2 \in \pi_2(\theta^+, \theta^-)$, we know that ϕ_2 is constrained to one of two choices. Each admits a holomorphic representative, but modulo 2 these cancel, so we see that $C = 0$. □

A minor modification of the previous argument yields the following:

Lemma 6.6. *Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram with $w, w' \in \mathbf{w}$ but $w'' \notin \mathbf{w}$. Let $(\Sigma, \boldsymbol{\alpha} \cup \alpha_0, \boldsymbol{\beta} \cup \beta_0, \mathbf{w} \cup \{w''\})$ be the diagram which is freely stabilized at w''. Suppose λ is a path from w to w'. If $A_{\overline{\lambda}}$ denotes the*

relative homology action on the diagram which is freely stabilized at w'', and A_λ is the map on the unstabilized diagram, then we have the following decomposition for sufficiently stretched almost complex structure:

$$A_{\overline{\lambda}} = \begin{pmatrix} A_\lambda & * \\ 0 & A_\lambda \end{pmatrix}.$$

The analogous statement holds if we replace the relative homology action with the absolute homology action A_γ for a closed curve γ which doesn't intersect w''. In particular we have that

$$S^+_{w''} A_\lambda = A_{\overline{\lambda}} S^+_{w''} \qquad \text{and} \qquad S^-_{w''} A_{\overline{\lambda}} = A_\lambda S^-_{w''}.$$

Similarly

$$S^+_{w''} A_\gamma = A_{\overline{\gamma}} S^+_{w''} \qquad \text{and} \qquad S^-_{w''} A_{\overline{\gamma}} = A_\gamma S^-_{w''}.$$

Proof. The proof follows as in Lemma 6.5, but multiplying the counts of disks by the appropriate $a(\lambda, \phi)$. □

6.3. **Dependence on gluing data.** This subsection is devoted to proving that maps S^+_w and S^-_w do not depend on the gluing data. Recall that in constructed the almost complex structure $\mathcal{J}(T)$ for given gluing data $\mathcal{J}$, we had to stretch ϕ to get a map ϕ_T. There is not a unique way to do this, though we have the following:

Lemma 6.7. *Suppose that $\mathcal{J} = (J_s, J'_s, B, B', r, r', p', \iota, \phi)$ is gluing data and that ϕ_T is as above. Let $\mathcal{J}(T)$ denote the almost complex structure defined above. Let $\mathcal{J}'(T)$ denote the almost complex structure constructed with a ϕ'_T instead of ϕ_T. Then we have that*

$$\Phi_{\mathcal{J}(T) \to \mathcal{J}'(T)} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

for any $T > 1$.

Proof. Given a fixed ϕ, the maps ϕ_T and ϕ'_T are determined up to isotopy which is constant on $A \cup A'$. As a result, there is an isotopy ψ_t of $\Sigma \# S^2$ which is constant outside of $\widetilde{A} \setminus (A \cup A')$ such that $\mathcal{J}'(T) = (\psi_1)_* \mathcal{J}(T)$. By naturality, the map $\Phi_{\mathcal{J}(T) \to \mathcal{J}'(T)}$ is chain homotopic to the diffeomorphism map $(\psi_1)_*$, which is the identity map on the obvious identification of the complexes $CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_\mathcal{J})$ and $CF^\circ_{\mathcal{J}'(T)}(\mathcal{H}_\mathcal{J})$ (as groups), i.e.

$$\Phi_{\mathcal{J}(T) \to \mathcal{J}'(T)} = (\psi_1)_* = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

□

We now consider the effect of increasing the neck length T.

Lemma 6.8. *Suppose that $\mathcal{J} = (J_s, J'_s, B, B', r, r', p', \iota, \phi)$ is a choice of gluing data. There is an N such that if $T, T' > N$ and $\mathcal{J}(T)$ and $\mathcal{J}(T')$ both achieve transversality, then*

$$\Phi_{\mathcal{J}(T) \to \mathcal{J}(T')} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. Consider a sequence of T_n and T'_n both approaching $+\infty$. For each n, pick a 1-parameter family of maps

$$\widehat{\phi}_{t,n} : \widetilde{A} \to S^1 \times \mathbb{R}$$

ranging over $t \in \mathbb{R}$, which interpolate ϕ_{T_n} and $\phi_{T'_n}$, i.e. for each fixed t, the map $\widehat{\phi}_{t,n}$ is conformal on A_0 and A_1 and for fixed $t \geq 1$ the map $\widehat{\phi}_{t,n}$ is a map from $\widetilde{A}$ to $S^1 \times [-a, T'_n + b]$, and for each $t \leq 1$ the map $\widehat{\phi}_{t,n}$ is a map from $\widetilde{A}$ to $S^1 \times [-a, T_n + b]$. Further assume that each of the complex structures $(\widehat{\phi}_{t,n})^*(S^1 \times \mathbb{R})$ on $\widetilde{A}$ has conformal length at least $a + b + \min(T_n, T'_n)$.

Let $\widehat{\mathcal{J}}_n$ denote the almost complex structure on $\Sigma \times I \times \mathbb{R}$ which is J_s on $(\Sigma \setminus r \cdot B)$, and is $\iota_* J'_s$ on $\iota(S^2 \setminus r' \cdot B')$. On $\widetilde{A}$, we define $\widehat{\mathcal{J}}_n$ to be $\mathrm{j}_t \times \mathrm{j}_\mathbb{D}$ where $\mathrm{j}_t = \widehat{\phi}^*_{t,n}(J_{S^1 \times \mathbb{R}})$.

As $n \to \infty$, the almost complex structures $\widehat{\mathcal{J}}_n$ approach the structure $J_s \vee J'_s$ on $(\Sigma \vee S^2) \times \mathbb{R} \times I$.

The map $\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')}$ can be computed by counting Maslov index 0 curves which are $\widehat{\mathcal{J}}_n$-holomorphic. Write

$$\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

We wish to show that $A = D = 1$ and $C = 0$ for sufficiently large n. First consider A and D. If $\phi = \phi_1 \# \phi_2$ is a Maslov index 0 homology disk in $\pi_2(\mathbf{x}\otimes\theta^+, \mathbf{y}\otimes\theta^+)$ or $\pi_2(\mathbf{x}\otimes\theta^-, \mathbf{y}\otimes\theta^-)$ in the diagram $\mathcal{H}_{\mathcal{J}}$, then using the combinatorial formula for the index, we have

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1.$$

By Lemma 6.3, we know that

$$\mu(\phi) = \mu(\phi_1) + 2m_2.$$

If ϕ has $\widehat{\mathcal{J}}_n$-holomorphic representatives for arbitrarily large n, then by passing to a Gromov limit of a subsequence, we know that ϕ_1 and ϕ_2 are represented by broken holomorphic disks. But then we know that $\mu(\phi_1) \geq 0$ and $\mu(\phi_2) \geq 0$ by transversality. Since $\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')}$ counts holomorphic disks with Maslov index zero, we conclude that $\mu(\phi_1) = m_2 = 0$ which implies that $\mu(\phi_2) = 0$ as well. Hence ϕ_1 has Maslov index zero, and since ϕ_1 has a J_s-holomorhic representative as a broken flowline, it must be a constant disk. By explicit examination of disks on the sphere, the only ϕ_2 in $\pi_2(\theta^+, \theta^+)$ or $\pi_2(\theta^-, \theta^-)$ with Maslov index zero are the constant disks. Hence we conclude that ϕ is a constant disk. Conversely the constant disks are always counted by $\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')}$ so we conclde that $A = D = 1$.

We now consider the component C. Suppose that $\phi = \phi_1 \# \phi_2$ is homology disk which has holomorphic representatives for arbitrarily large n. Again we have

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1.$$

Applying Lemma 6.3 to compute $\mu(\phi_2)$ shows that

$$\mu(\phi) = \mu(\phi_1) + 2m_2 + 1.$$

Since $m_2 \geq 0$, we conclude that $\mu(\phi_1) \leq -1$. If ϕ has holomorphic representatives for arbitrarily large n, then we know that ϕ_1 and ϕ_2 admit broken J_s-holomorphic representatives. In particular $\mu(\phi_1) \geq 0$, a contradiction. Hence we conclude that no disks are counted and $C = 0$. □

Definition 6.9. *We say that N is **sufficiently large** for gluing data $\mathcal{J}$, if $\Phi_{\mathcal{J}(T)\to\mathcal{J}(T')}$ is of the form in the previous lemma for all $T, T' \geq N$. We say T is **large enough to compute** $S_w^{\pm}$ for the gluing data $\mathcal{J}$ if $T > N$ for some N which is sufficiently large.*

The maps $S_w^{\pm}$ are now defined to be between the complexes $CF^{\circ}_{J_s}(\mathcal{H})$ and $CF^{\circ}_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}})$ for any T which is large enough to compute $S_w^{\pm}$ for $\mathcal{J}$. We now show that the maps $S_w^{\pm}$ are invariant under changing the map ϕ.

Lemma 6.10. *Suppose that choices of gluing data $\mathcal{J}$ and $\widetilde{\mathcal{J}}$ differ only in choice of the map ϕ. There is an N such that if $T > N$ and $\mathcal{J}(T)$ and $\widetilde{\mathcal{J}}(T)$ both achieve transversality, then we have*

$$\Phi_{\mathcal{J}(T)\to\widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. Suppose that $\phi : \widetilde{A} \to S^1 \times [-a, 1+b]$ and $\widetilde{\phi} : \widetilde{A} \to S^1 \times [-a, 1+b]$ are two choices of diffeomorphisms which are both conformal on A and A'. Suppose $T_n \to \infty$ and let ϕ_{T_n} and $\widetilde{\phi}_{T_n}$ be as above. We now construct a 1-parameter family of diffeomorphisms

$$\phi_{t,n} : \widetilde{A} \to [-a, T_n + b]$$

such that $\phi_{t,n} = \widetilde{\phi}_{T_n}$ for $t \geq 1$ and $\phi_{t,n} = \phi_{T_n}$ for $t \leq -1$ and such that $\phi_{t,n}$ is conformal on A and A' for all t. The existence of such a $\phi_{t,n}$ is seen as follows. The space $\mathrm{Diff}(S^1 \times I)$ is generated by Dehn twists around $S^1 \times \{\frac{1}{2}\}$, which are themselves isotopic to the identity map (since we are not required to fix the boundary). Hence there is an isotopy between ϕ_{T_n} and $\widetilde{\phi}_{T_n}$. Such an isotopy can clearly be taken to be only a rotation on $[-a, 0]$ and $[T, T + b]$. Since rotations are conformal maps on $S^1 \times \mathbb{R}$, we know that the isotopy between ϕ_{T_n} and $\widetilde{\phi}_{T_n}$ can be taken to be a conformal on A and A'. Hence such a $\phi_{t,n}$ exists. Here t denotes the $\mathbb{R}$ component of $\mathbb{D} = I \times \mathbb{R}$.

We now consider the almost complex structure $\widehat{\mathcal{J}}_n$ on $(\Sigma \# S^2) \times I \times \mathbb{R}$ defined to be J_s outside of $r \cdot B$ and J'_s outside of $r' \cdot B'$ and to be $\mathfrak{j}_t \times \mathfrak{j}_{\mathbb{D}}$ on $\widetilde{A} \times I \times \mathbb{R}$, where $\mathfrak{j}_t = \phi_{t,n}^*(J_{S^1 \times \mathbb{R}})$.

As we let $n \to \infty$, we get that $\widehat{\mathcal{J}}_n$ approaches the structure $J_s \vee J'_s$ on $(\Sigma \vee S^2) \times I \times \mathbb{R}$.

The map $\Phi_{\mathcal{J}(T_n) \to \widetilde{\mathcal{J}}(T_n)}$ can be computed by counting Maslov index zero $\widehat{\mathcal{J}}_n$-holomorphic disks on $\mathcal{H}_{\mathcal{J}}$. Write

$$\Phi_{\mathcal{J}(T_n) \to \widetilde{\mathcal{J}}(T_n)} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

An argument identical to the one in Lemma 6.8 shows that for all sufficiently large n, we have that $A = D = 1$ and $C = 0$, completing the claim. □

We now show that the maps $S_w^{\pm}$ are independent of the constants r and r' chosen in the gluing data.

Lemma 6.11. *Suppose that $\mathcal{J}$ and $\widetilde{\mathcal{J}}$ are two choices of gluing data which vary only in the choice of r and r'. Suppose also that ι and $\widetilde{\iota}$ are also related by restriction. Then there is an N such that if $T > N$ we have*

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. Notice that if $r < \widetilde{r}$, and $r' = \widetilde{r}'$, then the map ϕ_T can be translated in the $\mathbb{R}$ direction to make a map $\widetilde{\phi}_{T+c}$ for some $c > 0$. In particular we know that almost complex structure $\widetilde{\mathcal{J}}(T)$ is actually just equal to $\mathcal{J}(T + c)$ for some $c > 0$. Applying Lemma 6.8, then yields the desired form of $\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)}$. The same argument works if $\widetilde{r} < r$, and to change r'. □

We now show that the maps $S_w^{\pm}$ are independent of the almost complex structure J_s and the map ϕ used in the gluing data:

Lemma 6.12. *Suppose that $\mathcal{J} = (J_s, J'_s, B, B', r, r', p', \iota, \phi)$ is a choice of gluing data and $\widetilde{J}_s$ is another choice of almost complex structure on $\Sigma \times I \times \mathbb{R}$ which is split in a neighborhood containing B and such that ι is still an embedding of $\overline{S^2 \setminus r' \cdot B'} \hookrightarrow r \cdot B$. Then there is a $\widetilde{\phi}$ such that if $\widetilde{\mathcal{J}}$ denotes the gluing data $\widetilde{\mathcal{J}} = (J_s, \widetilde{J}'_s, B, B', r, r', p, \iota, \widetilde{\phi})$, then there is an N such that if $T > N$, we have an identification*

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} \Phi_{J_s \to \widetilde{J}_s} & * \\ 0 & \Phi_{J_s \to \widetilde{J}_s} \end{pmatrix}.$$

Proof. This proceeds roughly as before. We first construct an almost complex structure $\widehat{\mathcal{J}}(T)$ which can be used to count the transition map $\Phi_{\mathcal{J}(T) \to \mathcal{J}'(T)}$, and then we degenerate the almost complex structure by letting T approach $+\infty$, which will yield the total counts of disks via a Maslov index computation and a gluing argument.

First pick an interpolating conformal structure $\widehat{J}$ on $\Sigma \times I \times \mathbb{R}$ which agrees with J_s for $t \leq -1$ and with J'_s for $t \geq 1$. We can assume that $\widehat{J}$ is of the form $\mathfrak{j}_t \times \mathfrak{j}_{\mathbb{D}}$ for $t \in \mathbb{R}$ on a set containing $B \times I \times \mathbb{R}$. Pick a family $\psi_t : (B, w) \to (\mathbb{D}, 0)$ of conformal maps for the structure $\mathfrak{j}_t$. The map ψ_t is uniquely determined up to multiplication by a map $\omega : \mathbb{R} \to S^1$. A family of annuli $A_t \subseteq B$ for $t \in \mathbb{R}$ are defined by $B \setminus (r \cdot B)$ (notice that the set $r \cdot B$ depends on the conformal structure and hence varies with t). Using Lemma 6.11, we can assume that r is close enough to 1 so that $\overline{\iota(S^2 \setminus r' \cdot B')} \cap A_t = \varnothing$ for all t.

Let T_n be a sequence of numbers approaching $+\infty$. Let $\phi_{T_n} : \widetilde{A} \to S^1 \times [-a, T_n + b]$ be a choice of the stretched map ϕ.

Pick a 1-parameter family of smooth embeddings

$$\phi_{t,n} : \widetilde{A} \to S^1 \times \mathbb{R}$$

ranging over $t \in \mathbb{R}$, such that and if $t \leq -1$ then $\phi_{t,n} = \phi_{T_n}$ and such that $\phi_{t,n}$ maps A_t conformally to $[-a, 0]$ and maps A' conformally to $[T, T + b]$, and such that $\phi_{t,n}$ is constant in t for $t \geq 1$. Let $\widetilde{\phi}_{T_n}$ be $\phi_{t,n}$ for any $t \geq 1$. A map $\widetilde{\phi}$ can be constructed from $\widetilde{\phi}_{T_n}$ by scaling in the $\mathbb{R}$-component in the region $S^1 \times [0, T]$. We now define an almost complex structure $\widehat{\mathcal{J}}_n$ on $\Sigma \# S^2 \times I \times \mathbb{R}$ by defining it to be J_s outside on $\Sigma \setminus r \cdot B$, and and defining it to be J'_s on $S^2 \setminus r' \cdot B'$. On $\widetilde{A}$ it is $\mathfrak{j}_t \times \mathfrak{j}_{\mathbb{D}}$ where $\mathfrak{j}_t = \phi_{n,t}^*(J_{S^1 \times \mathbb{R}})$.

As $n \to \infty$, the almost complex structure $\widehat{\mathcal{J}}_n$ approaches the almost complex structure $\widehat{J} \vee J'_s$ on $\Sigma \vee S^2 \times I \times \mathbb{R}$.

Write

$$\Phi_{\mathcal{J}(T_n)\to\widetilde{\mathcal{J}}(T_n)} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

First consider A and D. Suppose that $\phi = \phi_1 \# \phi_2$ has Maslov index zero and has $\widehat{\mathcal{J}}_n$-holomorphic representatives for arbitrarily large n and that ϕ is in $\pi_2(\mathbf{x}\otimes\theta^+, \mathbf{y}\otimes\theta^+)$ or $\pi_2(\mathbf{x}\otimes\theta^-, \mathbf{y}\otimes\theta^-)$. Passing to a Gromov limit yields broken holomorphic strips U_1 and U_2 representing ϕ_1 and ϕ_2. By Lemma 6.3 we know that

$$0 = \mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 = \mu(\phi_1) + 2m_2.$$

Since ϕ_1 has holomorphic representatives we know $\mu(\phi_1) \geq 0$. By positivity of domains $m_2 \geq 0$. Hence $\mu(\phi_1) = 0$ and $m_2 = 0$. In principle the broken holomorphic strip U_1 could consist of many Maslov index zero strips, but only one strip can be $\widehat{J}$-holomorphic and the others must be J_s- or $\widetilde{J}_s$-holomorphic. Since generically there are no nonconstant J_s- or $\widetilde{J}_s$-holomorphic strips with Maslov index zero, we conclude that U_1 contains only a single nonconstant component u_1. There must be a matching component in U_2, namely a component u_2 with

$$\rho^{p_1}(u_1) = \rho^{p_2}(u_2).$$

We now claim that there are no other nonconstant components. Let ϕ_2' represent the homology class of u_2. Note that $m_1(u_1) = m_1(u_2)$ and so by positivity we know that $m_1(\phi_2 - \phi_2') = 0$. Since $\phi_2 - \phi_2'$ has holomorphic representatives as broken a broken holomorphic curve on $(S^2, \alpha_0, \beta_0, w, w')$, we conclude that $\mu(\phi_2 - \phi_2') \geq 0$ with equality iff $\mathcal{D}(\phi_2') = 0$. But then by an easy Maslov computation (the same one done in Lemma 6.5) we see that

$$0 = \mu(\phi) = \mu(\phi_1 \# \phi_2') + \mu(\phi_2 - \phi_2').$$

Since $\phi_1 \# \phi_2'$ has the holomorphic representative $u_1 \# u_2$, we conclude that $\mu(\phi_1 \# \phi_2') \geq 0$, and hence $\mu(\phi_2 - \phi_2') = 0$. Hence $\mathcal{D}(\phi_2 - \phi_2') = 0$ and the remaining components of U_2 are constant.

Conversely, given a $\widehat{J}$-holomorphic disk of Maslov index zero counted by $\Phi_{J_s\to\widetilde{J}_s}$, we can glue, as in Lemma 6.5 using Lemma 6.2. Hence we conclude that $A = D = \Phi_{J_s\to\widetilde{J}_s}$.

We finally consider component C. This is zero by an identical argument as the one in Lemma 6.8. □

We now consider changing the almost complex structure J_s' on S^2:

Lemma 6.13. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is a choice of gluing data and $\widetilde{J}_s'$ is another choice of almost complex structure on $S^2 \times I \times \mathbb{R}$ which is split in a neighborhood of B such that ι is still an embedding of $\overline{S^2 \setminus r' \cdot B'} \hookrightarrow r \cdot b$. Then there is a $\widetilde{\phi}$ such that if $\widetilde{\mathcal{J}}$ denotes the gluing data*

$$\widetilde{\mathcal{J}} = (J_s, J_s', B, B', r, r', p', \iota, \phi),$$

then there is an N such that if $T > N$, we have an identification o

$$\Phi_{\mathcal{J}(T)\to\widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. The argument is nearly identical to the proof of Lemma 6.12, so we omit it. □

We now show that the maps $S_w^{\pm}$ are invariant under the choice of embedding $\iota : \overline{S^2 \setminus r' \cdot B'} \hookrightarrow r \cdot B$.

Lemma 6.14. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is gluing data and suppose that $\widetilde{\iota}$ is another choice of embedding of $(\overline{S^2 \setminus r' \cdot B'}, w) \hookrightarrow (r \cdot B, w)$. Then there is a map $\widetilde{\phi}$ such that*

$$\widetilde{\mathcal{J}} = (J_s, J_s', B, B', r, r', p', \widetilde{\iota}, \widetilde{\phi})$$

is gluing data and

$$\Phi_{\mathcal{J}(T)\to\widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

for any $T > 1$.

Proof. The maps ι and $\widetilde{\iota}$ are isotopic relative $\mathbf{w} \cup \{w\}$. Let ψ_t be an isotopy of Σ fixing $\mathbf{w} \cup \{w\}$ such that $\psi_1 \circ \iota = \widetilde{\iota}$ and such that ψ_t is fixed outside of $r \cdot B$. We simply define $\widetilde{\phi} = \phi \circ \psi_1^{-1}$. We tautologically have $\widetilde{\mathcal{J}}(T) = (\psi_1)_* \mathcal{J}(T)$. By naturality of the Heegaard Floer complexes, we thus know that

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq (\psi_1)_* = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

□

We now consider the dependence on the choice of ball $B' \subseteq S^2$ and connected sum point $p' \in B'$:

Lemma 6.15. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is gluing data and that $\widetilde{B}'$ and $\widetilde{p}'$ are different choices of balls and connected sum point p'. Then there are choices $\widetilde{J}_s'$ of almost complex structure on S^2, embedding $\widetilde{\iota}$ and map $\widetilde{\phi}$ such that $\widetilde{\mathcal{J}} = (J_s, \widetilde{J}_s', B, \widetilde{B}', r, r', \widetilde{p}', \widetilde{\iota}, \widetilde{\phi})$ is gluing data and*

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

for all $T > 1$.

Proof. Let ψ_t be an isotopy of S^2 fixing $w \in S^2$, such that ψ_1 maps B' to $\widetilde{B}'$ and p' to $\widetilde{p}'$. We now set $\widetilde{J}_s' = (\psi_1)_* J_s$ and $\widetilde{\iota} = \iota \circ \psi_1^{-1}$. We can actually just set $\widetilde{\phi}$ equal to ϕ as functions. Tautologically have that $\mathcal{H}_{\mathcal{J}} = \mathcal{H}_{\widetilde{\mathcal{J}}}$ and $\mathcal{J}(T) = \widetilde{\mathcal{J}}(T)$, and hence $\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)}$ is tautologically the identity map. □

We now explore dependence on the final piece of a data, the ball $B \subseteq \Sigma$.

Lemma 6.16. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is gluing data and suppose that $\widetilde{B}$ is another ball in Σ containing w which doesn't intersect any of the $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves. Then there is gluing data*

$$\widetilde{\mathcal{J}} = (\widetilde{J}_s, J_s', \widetilde{B}, B', r, r', p', \widetilde{\iota}, \widetilde{\phi})$$

with

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Proof. We proceed similarly to the last two lemmas. Let ψ_t be an isotopy of Σ which fixes $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \{w\})$ and such that ψ_1 maps B to $\widetilde{B}$. We simply define $\widetilde{J}_s = (\psi_1)_* J_s$, $\widetilde{\iota} = \psi_1 \circ \iota$, and $\widetilde{\phi} = \phi \circ \psi_1^{-1}$. We have that $\widetilde{\mathcal{J}}(T) = (\psi_1)_* \mathcal{J}(T)$. By naturality we have that $\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq (\psi_1)_*$, which clearly has the desired form. □

6.4. **Free stabilization of the triangle maps.** In this section we prove a theorem about free stabilization and the triangle maps. The main consequence is that the free stabilization maps commute with maps which can be computed as a sequence of triangle maps, for instance the handleslide maps or isotopy maps. The result has some similarities with the quasi-stabilization argument in [MO10], though our argument is somewhat different.

Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a strongly $\mathfrak{s}$-admissible Heegaard triple for $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma}})$ and suppose that $w \in \Sigma \setminus (\mathbf{w} \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma})$ is a new basepoint. Fix curves α_0, β_0 and γ_0 which are pairwise isotopic and each bound a small disk containing w and which don't intersect any other $\boldsymbol{\alpha}, \boldsymbol{\beta}$ or $\boldsymbol{\gamma}$ curves. Further assume that $|\sigma_0 \cap \tau_0| = 2$ for distinct $\sigma, \tau \in \{\alpha, \beta, \gamma\}$. Define

$$\overline{\boldsymbol{\alpha}} = \boldsymbol{\alpha} \cup \{\alpha_0\}$$

and define $\overline{\boldsymbol{\beta}}$ and $\overline{\boldsymbol{\gamma}}$ similarly. Define

$$\overline{\mathcal{T}} = (\Sigma, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}}, \mathbf{w} \cup \{w\}).$$

In light of Lemma 4.11, a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma}})$ uniquely determines a Spin^c structure $\bar{\mathfrak{s}} \in \mathrm{Spin}^c(X_{\overline{\boldsymbol{\alpha}}\overline{\boldsymbol{\beta}}\overline{\boldsymbol{\gamma}}})$, An easy argument shows that $\mathcal{T}$ is strongly $\mathfrak{s}$-admissible iff $\overline{\mathcal{T}}$ is strongly $\bar{\mathfrak{s}}$ admissible. The diagram $\overline{\mathcal{T}}$ is shown in Figure 6.

The main result of this section is the following:

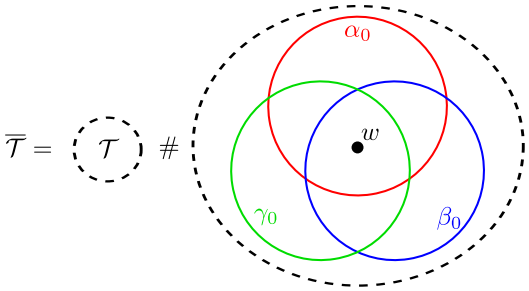

FIGURE 6. The Heegaard triple $\overline{\mathcal{T}}$.

Theorem 6.17. *If $\mathcal{J}$ is gluing data, as in the previous section, then for all sufficiently large T such that $\mathcal{J}(T)$ achieves transversality, we have that*

$$F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^+)) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix},$$

where the matrix on the right is an expansion of the map on the left in the variable $\cdot$, which can be either x^+ or x^-.

Similarly, for sufficiently stretched almost complex structure, we have

$$F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}((\mathbf{x}\times x^+)\otimes(\mathbf{y}\times\cdot)) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix}.$$

As with the previous neck stretching arguments, we need a lemma about counts of disks in $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$ satisfying a certain matching condition. This is similar to [MO10, Lem. 5.16], which is itself very similar to [OS08, Lem. 6.4]. The first step is a Maslov index computation. Consider the Heegaard triple $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$ shown in Figure 7. We first perform a Maslov index computation for triangles on $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$:

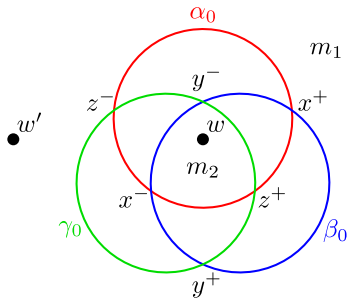

FIGURE 7. The Heegaard triple $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$ with $x^\pm,y^\pm,z^\pm$ and multiplicities m_1 and m_2 labelled.

Lemma 6.18. *Suppose $\psi\in\pi_2(a,b,c)$ is a triangle on $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$. If $\psi\in\pi_2(x^+,y^+,z^+),\pi_2(x^-,y^+,z^-)$, or $\pi_2(x^+,y^-,z^-)$. Then*

$$\mu(\psi) = 2m_1+2m_2.$$

If $\psi\in\pi_2(x^+,y^+,z^-)$ then

$$\mu(\psi) = 2m_1+2m_2+1.$$

If $\psi\in\pi_2(x^-,y^+,z^+)$ or $\pi_2(x^+,y^-,z^+)$, then

$$\mu(\psi) = 2m_1+2m_2-1.$$

Proof. In all cases, the formula can be easily verified for a particular "simple" homology class of triangles. The formula in fully generality then follow since both respect splicing in homology classes of boundary degenerations, and any two homology triangles with the same vertices differ by a sum of domains corresponding to boundary degenerations. □

Let $p \in S^2$ be a point in the same region as w'. Let $x \in \alpha_0 \cap \beta_0, y \in \beta_0 \cap \gamma_0$ and $z \in \alpha_0 \cap \gamma_0$. Given a divisor $\mathbf{d} \in \mathrm{Sym}^k(\Delta)$, and a homology triangle $\psi \in \pi_2(x, y, z)$ we define

$$\mathcal{M}(\psi, \mathbf{d}) = \{u \in \mathcal{M}(\psi) : \rho^p(u) = \mathbf{d}\}$$

and

$$M_{(x,y,z)}(\mathbf{d}) = \sum_{\substack{\psi \in \pi_2(x,y,z) \\ n_w(\psi)=0 \\ \mu(\psi)=2k}} \#\mathcal{M}(\psi, \mathbf{d}).$$

Lemma 6.19. *If (x, y, z) is one of the tuples $(x^+, y^+, z^+), (x^-, y^+, z^-)$ or (x^+, y^-, z^-), then the number $M_{(x,y,z)}(\mathbf{d})$ is independent of $\mathbf{d}$ for generic $\mathbf{d}$ and in fact*

$$M_{(x,y,z)}(\mathbf{d}) \equiv 1 (\mathrm{mod}\, 2).$$

Proof. The proof is nearly identical to the proof of [OS08, Lem. 6.4]. We first show that $M_{(x,y,z)}(\mathbf{d})$ is independent of $\mathbf{d}$ using a Gromov compactness argument. Suppose that $\mathbf{d}_t$ is a path in $\mathrm{Sym}^k(\Delta)$ between $\mathbf{d}_0$ and $\mathbf{d}_1$. We then consider the 1-manifold

$$\mathcal{M}_{(x,y,z)}(\mathbf{d}_t) = \bigcup_{\substack{t \in I \\ \psi \in \pi_2(x,y,z) \\ n_w(\psi)=0 \\ \mu(\psi)=2k}} \mathcal{M}(\psi, \mathbf{d}_t).$$

We count the ends of $\mathcal{M}_{(x,y,z)}(\mathbf{d}_t)$. There are several types of ends. The first type corresponds to setting $t = 0$ or $t = 1$, namely

$$\bigcup_{\substack{\psi \in \pi_2(x,y,z) \\ n_w(\psi)=0 \\ \mu(\psi)=2k}} \mathcal{M}(\psi, \mathbf{d}_i).$$

The other types of ends correspond to homology classes of triangles ψ degenerating into broken holomorphic triangles. Let x', y' and z' denote the "other" generators of $\alpha_0 \cap \beta_0$, $\beta_0 \cap \gamma_0$ and $\alpha_0 \cap \gamma_0$ respectively, i.e. the generators in the opposite grading as x, y and z. In principle a sequence of holomorphic strips of index $2k$ could break into a collection of many holomorphic strips, boundary degenerations, and closed surfaces, whose total homology represented the original strip. In the ends of $\mathcal{M}_{(x,y,z)}(\mathbf{d}_t)$, one of the components must satisfy the matching condition with a divisor $\mathbf{d}_t$, and hence using the Maslov index computation in Lemma 6.18, the possible degenerations are limited to the following types:

(a)
$$\bigcup_{\substack{t \in I \\ \phi \in \pi_2(x,x') \\ \psi \in \pi_2(x',y,z) \\ n_w(\phi)=n_w(\psi)=n_{w'}(\phi)=0 \\ \mu(\phi)=1, \mu(\psi)=2k-1}} \mathcal{M}(\psi, \mathbf{d}_t) \times \widehat{\mathcal{M}}(\phi),$$

(b)
$$\bigcup_{\substack{t \in I \\ \phi \in \pi_2(y,y') \\ \psi \in \pi_2(x,y',z) \\ n_w(\phi)=n_w(\psi)=n_{w'}(\phi)=0 \\ \mu(\phi)=1, \mu(\psi)=2k-1 \\ n_{w'}(\phi)=0}} \mathcal{M}(\psi, \mathbf{d}_t) \times \widehat{\mathcal{M}}(\phi),$$

and

(c) $$\bigcup_{\substack{t\in I\\ \phi\in\pi_2(z',z)\\ \psi\in\pi_2(x,y,z')\\ n_w(\phi)=n_w(\psi)=n_{w'}(\phi)=0\\ \mu(\phi)=1,\mu(\psi)=2k-1}} \mathcal{M}(\psi,\mathbf{d}_t)\times\widehat{\mathcal{M}}(\phi).$$

The total number of elements in any of the ends of type (a),(b) or (c) is even, since the differential of the hat complexes $\widehat{CF}(S^2,\sigma,\tau,w,w')$ vanishes for any two distinct $\sigma,\tau\in\{\alpha_0,\beta_0,\gamma_0\}$.

Hence the total number of ends, modulo 2, is just

$$\#\mathcal{M}(\mathbf{d}_1)-\#\mathcal{M}(\mathbf{d}_0),$$

and hence that quantity is zero.

We now wish to show that $M(\mathbf{d})\equiv 1$, modulo 2. Let $\mathbf{d}_T$ consist of k points spaced at least T apart in Δ such that as $T\to\infty$ the points approach e_β. As $T\to\infty$, the set of curves splits into k Maslov index 2 β-boundary degenerations. There is a unique Maslov index 2 homotopy class of β boundary degeneration ϕ with $n_w(\phi)=0$, namely the domain which has multiplicity one in the complement of the region in $S^2\setminus\beta_0$ which doesn't contain w, and zero elsewhere. By [OS08, Thm. 5.5], we know that a boundary degeneration of Maslov index 2 has a unique holomorphic representative modulo the action of $PSL_2(\mathbb{R})$. Hence we conclude that $\lim_{T\to\infty}\#\mathcal{M}(\psi,\mathbf{d}_T)=1(\mathrm{mod}\,2)$ by gluing, from which the claim follows.

□

We are now in position to prove the main result of this subsection:

Proof of Theorem 6.17. We will focus on computing the map $F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^+))$ in terms of $F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y})$, computing the second map in the theorem statement follows an identical argument.

Write

$$F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^+))=\begin{pmatrix}A(\mathbf{x},\mathbf{y}) & B(\mathbf{x},\mathbf{y})\\ C(\mathbf{x},\mathbf{y}) & D(\mathbf{x},\mathbf{y})\end{pmatrix}.$$

We will start by computing A and D. Suppose that $\psi=\psi_1\#\psi_2$ is a homology triangle in $\pi_2(\mathbf{x}\times x^+,\mathbf{y}\times y^+,\mathbf{z}\times z^+)$ or $\pi_2(\mathbf{x}\times x^-,\mathbf{y}\times y^+,\mathbf{z}\times z^-)$ of Maslov index zero. Using the combinatorial formula for the Maslov index, we have that

$$0=\mu(\psi_1\#\psi_2)=\mu(\psi_1)+\mu(\psi_2)-2m_1.$$

Using Lemma 6.18, we have

$$\mu(\psi)=\mu(\psi_1)+2m_2.$$

If ψ has $\mathcal{J}(T)$-holomorphic representatives u_T for arbitrarily large T, then a Gromov compactness argument shows that ψ_1 and ψ_2 have representatives as broken holomorphic triangles, U_1 and U_2 respectively. In particular $\mu(\psi_1)\geq 0$. In light of the equality $0=\mu(\psi_1)+2m_2$, from positivity of domains we conclude that $\mu(\psi_1)=2m_2=0$. The collection U_1 is a broken holomorphic triangle, which means that it is a collection of a single holomorphic triangle, a collection of holomorphic strips, boundary degenerations, and closed surfaces. Since the total Maslov index is zero, we conclude that the only nonconstant component of U_1 is a single holomorphic triangle, u_1. There must be a component of U_2, say u_2, with the matching condition

$$\rho^{p_1}(u_1)=\rho^{p_2}(u_2).$$

Say that u_2 is represented by the homology triangle ψ_2'. Then $\psi_2-\psi_2'$ admits a representative as a broken holomorphic curve, and hence $\mu(\psi_2')\geq 0$, with equality iff $\mathcal{D}(\psi_2')=0$, since ψ_2' is represented by holomorphic strips, boundary degenerations, and closed surfaces.

Since $\psi_1\#\psi_2'$ has a holomorphic representative, $u_1\#u_2$, we conclude that $\mu(\psi_1\#\psi_2')\geq 0$. Because of the matching condition,we have $m_1(\psi_2')=m_1(\psi_1)=m_1(\psi_2)$ and hence

$$m_1(\psi_2-\psi_2')=0.$$

Hence we have that

$$0=\mu(\psi_1\#\psi_2)=\mu(\psi_1\#\psi_2')+\mu(\psi_2-\psi_2').$$

Combining this with our previous observations, we conclude that $\mu(\psi_1\#\psi_2') = 0$ and $\mu(\psi_2 - \psi_2') = 0$ and hence $\psi_2' = \psi_2$. Hence the only nonconstant component of U_2 is u_2. We thus conclude that the pair (u_1, u_2) were a preglued pair.

Conversely, given a preglued pair, we can glue to form a genuine holomorphic triangle for sufficiently stretched almost complex structure. By Lemma 6.19, given a holomorphic triangle u_1 repesenting a $\psi_1 \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$, among holomorphic triangles representing classes in $\pi_2(x^+, y^+, z^+)$, there is one (modulo 2), holomorphic triangle satisfying the matching condition with u_1. Hence for sufficiently stretched almost complex structures, we have an identification of

$$A(\mathbf{x}, \mathbf{y}) = F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}).$$

The same reasoning shows that we have an identification

$$D(\mathbf{x}, \mathbf{y}) = F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}).$$

We finally wish to show that $B(\mathbf{x}, \mathbf{y}) = 0$. Suppose that $\psi = \psi_1\#\psi_2 \in \pi_2(\mathbf{x} \times x^+, \mathbf{y} \times y^+, \mathbf{z} \times z^-)$ is a triangle with Maslov index zero. By 6.18 we have that

$$0 = \mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1) + 2m_2 + 1.$$

If ψ has $\mathcal{J}(T)$-holomorphic representatives for arbitrarily large T, then we conclude that ψ_1 must have a representative as a broken holomorphic triangle, and hence $\mu(\psi_1) \geq 0$. Since $m_2 \geq 0$ by positivity, we conclude that

$$\mu(\psi_1) + 2m_2 + 1 \geq 1,$$

and hence no triangles are counted for sufficiently stretched almost complex structure, so $B(\mathbf{x}, \mathbf{y}) = 0$. □

6.5. **Topological preliminaries about based embeddings of Heegaard surfaces.** In this section we prove a lemma about isotopies of Heegaard surfaces which will be useful for proving naturality of the free stabilization maps, as well as the 1- and 3-handle maps. We begin with a two parameter isotopy extension theorem. We leave the proof to the reader, since it is elementary:

Lemma 6.20. *Let M be a closed manifold and let $\Phi_t : M \to M$ be an isotopy. Let N be an embedded submanifold and let $\phi_t = \Phi_t|_N$. Suppose that ψ_t is another isotopy of N such that $\psi_0 = \phi_0$ and $\psi_1 = \phi_1 = \mathrm{id}$ and suppose there is a homotopy $(h_t)_s$ through isotopies between ϕ_t and ψ_t such that $(h_0)_s$ and $(h_1)_s$ are constant in s. Then there is another isotopy $\Psi_t : M \to M$ such that $\Psi_0 = \Phi_0 = \mathrm{id}_M$ and $\Psi_1 = \Phi_1$ and $\Psi_t|_N = \psi_t$. Furthemore the isotopy Ψ_t can be taken to be equal to Φ_t outside of a given open neighborhood of $h(N \times I \times I) \subseteq M$.*

Lemma 6.21. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ and $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w})$ are Heegaard diagrams for $(Y, \mathbf{w})$ and that $\mathbf{p} \subseteq Y$ is a finite collection of points such that $\mathbf{p} \subseteq \Sigma \cap \Sigma'$. Suppose further that $\mathcal{H}$ and $\mathcal{H}'$ are related by an isotopy of Y fixing $\mathbf{w}$. Then $\mathcal{H}$ and $\mathcal{H}'$ are related by an isotopy F_t fixing $\mathbf{w}$ such that $\mathbf{p} \subseteq F_t(\Sigma)$ for all t.*

Proof. Let $G_t : Y \to Y$ be an isotopy fixing $\mathbf{w}$ which maps $\mathcal{H}$ to $\mathcal{H}'$. Note that $\mathbf{p} \subseteq G_t(\Sigma)$ for all t iff $G_t^{-1}(\mathbf{p}) \subseteq \Sigma$ for all t. Let $\Phi_t = G_t^{-1}$ and let $\phi_t = \Phi_t|_{\mathbf{p}}$. Note that ϕ_t is a collection of paths starting at $\mathbf{p}$ and ending at a collection of points $\mathbf{q} \subseteq \Sigma$.

Since $\pi_1(\Sigma)$ surjects onto $\pi_1(Y)$, we can homotope ϕ_t with a homotopy $(h_t)_s$ so that $\mathrm{im}\,\phi_t \subseteq \Sigma$. We can assume that the image of $(h_t)_s$ avoids the basepoints $\mathbf{w}$. Applying Lemma 6.20, there is an isotopy Ψ_t of Y such that $\Psi_t(\mathbf{p}) \subseteq \Sigma$ for all t, and such that $\Psi_0 = \Phi_0$ and $\Psi_1 = \Phi_1$. Furthermore, since the image of $(h_t)_s$ avoids $\mathbf{w}$, we can assume that Ψ_t fixes $\mathbf{w}$, since Φ_t also fixes $\mathbf{w}$. Set $F_t = \Psi_t^{-1}$. Since $\Psi_1 = \Phi_1$, we know that F_t is an isotopy between $\mathcal{H}$ and $\mathcal{H}'$. By construction, $F_t^{-1}(\mathbf{p}) \subseteq \Sigma$. □

6.6. **Invariance of free stabilization maps.** In this subsection we combine the work of the previous subsections to show that the free stabilization maps yield well defined maps on Heegaard Floer homology.

Theorem 6.22. *The free stabilization maps $S_w^{\pm}$ described above yield well defined maps*

$$S_w^+ : HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S_w^- : HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

on homology.

Proof. Suppose that $\mathcal{H}_1 = (\Sigma_1, \boldsymbol{\alpha}_1, \boldsymbol{\beta}_1, \mathbf{w})$ and $\mathcal{H}_2 = (\Sigma_2, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \mathbf{w})$ are two Heegaard diagrams for $(Y, \mathbf{w})$ which are subordinate to $w \in Y \setminus \mathbf{w}$, and that $\mathcal{J}_i = (J_{s,i}, J'_{s,i}, B_i, B'_i, r_i, r'_i, p', \iota_i, \phi_i)$ are choices of almost complex structures for $\mathcal{H}_i$. We wish to show that that the following diagram commutes up to chain homotopy:

$$\begin{array}{ccc} CF^\circ_{J_{s,1}}(\mathcal{H}_1) & \xrightarrow{S^+_w} & CF^\circ_{\mathcal{J}_1(T_1)}((\mathcal{H}_1)_{\mathcal{J}_1}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}_1 \to \mathcal{H}_2}} & & \downarrow{\scriptstyle \Phi_{(\mathcal{H}_1)_{\mathcal{J}_1} \to (\mathcal{H}_2)_{\mathcal{J}_2}}} \\ CF^\circ_{J_{s,2}}(\mathcal{H}_2) & \xrightarrow{S^+_w} & CF^\circ_{\mathcal{J}_2(T_2)}((\mathcal{H}_2)_{\mathcal{J}_2}). \end{array}$$

for any T_i which are large enough to compute $S^\pm_w$ for the gluing data $\mathcal{J}_i$. According to Lemma 4.20, the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ can be written as a composition of moves corresponding to four moves: change of almost complex structure, triangle maps (isotopies and handleslides), (1,2)-stabilization and destabilization, and isotopies of Σ inside of Y fixing $\mathbf{w}$.

We first address changing the gluing data $\mathcal{J}$ for a fixed Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$. Combining Lemmas 6.7, 6.8, 6.10, 6.11, 6.12, 6.13, 6.14, 6.15, and 6.16, we see that the change of almost complex structure map

$$\Phi_{\mathcal{J}_1(T_1) \to \mathcal{J}_2(T_2)} \simeq \begin{pmatrix} \Phi_{J_{s,1} \to J_{s,2}} & * \\ 0 & \Phi_{J_{s,1} \to J_{s,2}} \end{pmatrix},$$

for T_i which are large enough to compute $S^\pm_w$ for gluing data $\mathcal{J}_i$, which immediately implies that S^+_w and S^-_w commutes with the change of almost complex structure map. We leave it to the reader to combine these lemmas to produce the above matrix since it is quite easy, though we do make a comment about the definition of T being "large enough to compute $S^\pm_w$". In going between between the two almost complex structures $\mathcal{J}_1(T_1)$ and $\mathcal{J}_2(T_2)$, we need to go through intermediate almost complex structures $\mathcal{J}_i(T)$ for T which is large enough to compute $S^\pm_w$ for both $\mathcal{J}_1$ and $\mathcal{J}_2$ (and potentially even larger, so that we know the forms of the intermediate transition maps). To do this we need Lemma 6.8 to show that as we increase T, eventually the change of almost complex structure maps become predictable.

We now consider maps which can be computed with a triangle map (handleslide maps and isotopy maps). By Theorem 6.17, the free stabilization maps commute with these triangles maps. Upon direct examination, the (1,2)-stabilization maps commute with the free stabilization maps.

We finally consider maps induced by isotopies of Y relative the points $\mathbf{w}$. Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram with $w \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ and suppose $\phi_t : Y \to Y$ is an isotopy which fixes $\mathbf{w}$ and such that $w \in \phi_1(\Sigma)$. Letting $\phi = \phi_1$, we need to show that the following diagram commutes up to chain homotopy for some choices of gluing data $\mathcal{J}_1$ on $\mathcal{H}$ and $\mathcal{J}_2$ on $\phi_*\mathcal{H}$:

$$\begin{array}{ccc} CF_{J_{s,1}}(\mathcal{H}) & \xrightarrow{S^+_w} & CF^\circ_{\mathcal{J}_1(T_1)}(\mathcal{H}_{\mathcal{J}_1}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H} \to \phi_*\mathcal{H}}} & & \downarrow{\scriptstyle \Phi_{\mathcal{H}_{\mathcal{J}_1} \to (\phi_*\mathcal{H})_{\mathcal{J}_2}}} \\ CF_{J_{s,2}}(\phi_*\mathcal{H}) & \xrightarrow{S^+_w} & CF^\circ_{\mathcal{J}_2(T_2)}((\phi_*\mathcal{H})_{\mathcal{J}_2}) \end{array} \tag{6}$$

To this end, by Lemma 6.21, we can assume that $w \in \phi_t(\Sigma)$ for all t. Let $w_t = \phi_t^{-1}(w)$. The map w_t can be viewed as a smooth isotopy of the point w in Σ. Extend the isotopy w_t to an isotopy of $\Sigma \times I \to \Sigma$ which fixes $\mathbf{w}$, and then extend this to an isotopy $W_t : Y \times I \to Y$ which fixes $\mathbf{w}$ and maps Σ to Σ. Consider the two isotopies

$$\rho_t = W_t$$

and

$$\psi_t = \phi_t \circ W_t^{-1}.$$

Note that

$$\psi_1 \circ \rho_1 = \phi_1.$$

Also note that ψ_t fixes $\mathbf{w} \cup \{w\}$, and ρ_t fixes $\mathbf{w}$ and maps Σ diffeomorphically to Σ. Hence to show that the diagram in (6) commutes for an arbitrary isotopy ϕ_t fixing $\mathbf{w}$, it is sufficient to show that the diagram in (6) commutes if ϕ_t either

(1) fixes $\mathbf{w} \cup \{w\}$;
(2) fixes $\mathbf{w}$ and maps Σ into Σ.

Note that if ϕ_t fixes $\mathbf{w} \cup \{w\}$, then the commutativity of the diagram in (6) is tautological because we can take $\mathcal{J}_2 = \phi_* \mathcal{J}_1$ and $\Phi_{\mathcal{H}_{\mathcal{J}_1} \to (\phi_* \mathcal{H})_{\mathcal{J}_2}} = (\phi_1)_*$.

Now suppose that ϕ_t is an isotopy which fixes $\mathbf{w}$ but which maps Σ to Σ. The map $\Phi_{\mathcal{H} \to (\phi_* \mathcal{H})}$ can be computed simply as the tautological map $(\phi_1)_*$, but by naturality (cf. [JT12, Thm. 2.39]), it can also be computed as a composition of isotopies of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves, and change of almost complex structure maps. For our present purpose, we can thus assume that the underlying Heegaard surfaces for $\mathcal{H}$ and $\phi_* \mathcal{H}$ are the same, which allows us to pick the gluing data $\mathcal{J}_1$ and $\mathcal{J}_2$ to be equal. The maps corresponding to isotopies of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves can be computed as compositions of triangle maps. To such a move on $\mathcal{H}$, there is a corresponding move on $\mathcal{H}_{\mathcal{J}}$, which can also be computed as a triangle map. By Theorem 6.17 we know the free stabilization maps commute with such moves with the triangle maps. Hence we know that the diagram in (6) commutes up to chain homotopy for some choice of gluing data, completing the proof. □

7. Graph action map

In this section we describe a sort of 3-dimensional graph TQFT as an intermediate step to defining the graph cobordisms. Roughly speaking, to a graph $\Gamma \subseteq Y$ with two distinguished disjoint subsets of vertices, V_0 and V_1, we will construct a map from $HF^\circ(Y_1, V_0, \mathfrak{s})$ to $HF^\circ(Y_2, V_1, \mathfrak{s})$, called the graph action map. Due to some technical requirements, we will have to construct it in a slightly larger generality, allowing additional vertices $\mathbf{w}_0$ which don't interact with the graph Γ. We now precisely define the graphs involved:

Definition 7.1. *A **ribbon graph** Γ is a graph equipped with a cyclic ordering of the edges incident to each vertex. A **strong ribbon graph** is a graph Γ equipped with an absolute ordering on the vertices indicent to each vertex.*

Definition 7.2. *Suppose that Y^3 is a closed 3-manifold. We call a tuple $\mathcal{G} = (\Gamma, V_0, V_1, (\sigma, \mathfrak{P}), \mathbf{w}_0)$ a strong ribbon (resp. ribbon) **flow graph** if*

(1) Γ is an embedded, strong ribbon (resp. ribbon) graph;
(2) $(\sigma, \mathfrak{P})$ is a coloring of Γ and the basepoints $\mathbf{w}_0$;
(3) $\mathbf{w}_0$ is a collection of basepoints in Y with $\mathbf{w}_0 \cap \Gamma = \varnothing$;
(4) $V_0, V_1 \subseteq V(\Gamma)$ are nonempty subsets of the vertices $V(\Gamma)$ with $V_0 \cap V_1 = \varnothing$;
(5) the inclusions $\pi_0(V_0 \cup \mathbf{w}_0) \to \pi_0(Y)$ and $\pi_0(V_1 \cup \mathbf{w}_0) \to \pi_0(Y)$ are both surjective.
(6) each vertex in V_0 or V_1 has valence 1;
(7) there are no vertices of $V(\Gamma)$ which are not incident to an edge.

The main theorem of this section is the following:

Theorem B. *Suppose that $(\mathcal{G}, \Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is a strong ribbon flow graph in Y and $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. If $V \subseteq V(\Gamma)$ is a collection of vertices with $V_0 \subseteq V$, then then there is a well defined map*

$$A^V_{\mathcal{G},\mathfrak{s}} : HF^\circ(Y, V \cup \mathbf{w}_0, \sigma|_{V \cup \mathbf{w}_0}, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, V_1 \cup \mathbf{w}_0, \sigma|_{V_1} \cup \mathbf{w}_0, \mathfrak{P}, \mathfrak{s}).$$

The map $A^V_{\mathcal{G},\mathfrak{s}}$ depends only on the ordering of the edges incident to a vertex $v \notin V$ up to cyclic permutation. The map $A^V_{\mathcal{G},\mathfrak{s}}$ is functorial in the following sense. Suppose that $\mathcal{G} = (\Gamma, V_0, V_2, \sigma, \mathbf{w}_0)$ is a flow graph which decomposes as the composition of two flow graphs $\mathcal{G} = \mathcal{G}_2 \circ \mathcal{G}_1$ where $\mathcal{G}_i = (\Gamma_i, V_{i-1}, V_i, \sigma_i)$. Then

$$A^{V_1}_{\mathcal{G}_2,\mathfrak{s}} \circ A^{V_0}_{\mathcal{G}_1,\mathfrak{s}} = A^{V_0}_{\mathcal{G},\mathfrak{s}}.$$

Of course the zero maps satisfy the above theorem, but our maps will be much more interesting than the zero maps. To define the graph action maps, we will break a graph into small pieces, define the graph action map for each piece, and then compose all of the maps together.

Definition 7.3. *We say that a flow graph $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is **elementary** if it has one of the following forms:*

(1) $|V_1| = |V_0|$ and each edge of Γ connects a vertex of V_0 to a vertex of V_1;
(2) There is a single vertex v of Γ which is not in $V_0 \cup V_1$ and all edges of Γ connect either V_0 to V_1 or to a point of $V_0 \cup V_1$ to v;

(3) $|V_1| = |V_0| \pm 2$*, all edges except for one connect a vertex in* V_0 *to a vertex in* V_1 *and there is a single edge which connects two vertices in* V_0 *to each other, or two vertices in* V_1 *to each other. We divide elementary flow graphs of type* (3) *into two subtypes: type* $(3i)$ *corresponds to the edge connecting two vertices in* V_0*, and type* $(3ii)$ *corresponds to the edge connecting two vertices in* V_1*.*

Examples of the three types of elementary flow graphs are depicted in Figure 8 .

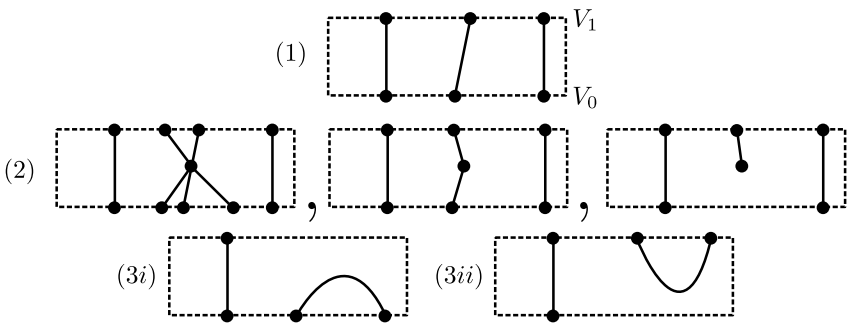

FIGURE 8. Examples of the three types of elementary flow graphs. We view these graphs as flowing from bottom to top. We divide type (3) into two subtypes: type $(3i)$ and $(3ii)$.

Definition 7.4. *Given a flow graph* $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ *with a distinguished set of vertices* $V \subseteq V(\Gamma)$ *with* $V_0 \subseteq V$*, we define a* ***Cerf decomposition*** *of a flow graph* $(\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ *to be a collection of flow graphs* $\mathcal{G}_i = (\Gamma_i, V_{0,i}, V_{1,i}, \sigma_i, \mathbf{w}_0 \cup (V \setminus V(\Gamma_i)))$ *for* $i = 0, \dots n$ *such that*

(1) $V_0 = V_{0,0}$ *and* $V_1 = V_{1,n}$*;*
(2) each $\mathcal{G}_i$ *is elementary;*
(3) $V_{0,i} = V_{1,i-1}$ *for* $i > 0$*;*
(4) $\cup_i \Gamma_i$ *is a subdivision of* Γ*;*
(5) the $V_{0,i}$ *and* $V_{1,i}$ *do not contain any vertices of* Γ *(the original, nonsubdivided graph) except for* $V_{0,0}$ *and* $V_{1,n}$*;*
(6) if $\mathcal{G}_i$ *is an elementary flow graph of type (2), then the middle vertex,* v*, is a vertex in the nonsubdivided graph* Γ*.*
(7) $\sigma_i = \sigma|_{\Gamma_i \cup \mathbf{w}_0}$*.*

We emphasize the distinction between the original vertices of Γ, and the vertices of the Cerf decomposition. We refer to the graph Γ with the original collection of vertices as the nonsubdivided graph, and we refer to the vertices of the Cerf decomposition as the subdivided graph. A vertex of Γ which is not in V_0 or V_1 is never an incoming or outgoing vertex for any of the elementary flow graphs in a Cerf decomposition. Instead a vertex of Γ which is not in V_0 or V_1 is always the middle vertex of an elementary flow graph of type 2.

A Cerf decomposition can be obtained by picking a Morse function $f : \Gamma \to [0, 1]$ with $f^{-1}(i) = V_i$ for $i = 0, 1$ which has connected level sets. Each edge e can be viewed as the image of an immtesion of $e : I \to Y$, which is embedded except possibly at the endpoints, where $e(0)$ and $e(1)$ might coincide. By a Morse function $f : \Gamma \to [0, 1]$, we mean a function f such that $f \circ e$ is Morse for each edge $e : I \to Y$. A Cerf decomposition then results by picking level sets of a generic Morse function f such that between any two consecutive level sets, there is at most 1 vertex of Γ or critical point of f. Conversely any Cerf decomposition arises from a Morse function on Γ in this way, since given a Cerf decomopsition of a flow graph, a Morse function can easily be built on each elementary flow graph, and glued together to form a Morse function on the entire flow graph which yields the original Cerf decomposition.

Moving between two Cerf decompositions corresponds to doing Cerf theory on compact 1-manifolds with boundary. With the above in mind, we have the following:

Lemma 7.5. *Any two Cerf decompositions of a flow graph can be connected by moves corresponding to the following set of moves and their inverses on level of Morse functions:*

(1) *inserting a level set, splitting an elementary flow graph of type* (n) *(for* $n \in \{1, 2, 3i, 3ii\}$*), with two elementary flow graphs, one of type* (n) *and one of type (1);*
(2) *performing a birth-death move for a pair of critical points along the interior of an edge, replacing two elementary flow graphs of type* $(3i)$ *and* $(3ii)$ *respectively with one of type (1);*
(3) *eliminating a critical point along an edge by sending it to a vertex;*
(4) *exchanging the relative values of two critical points on different edges of the subdivided graph which don't share a common vertex;*
(5) *exchanging the relative values of a vertex* v *of* Γ *(the unsubdivided graph) and a critical point which is on an edge such that there is not an edge (of the subdivided graph) which is incident to both* v *and contains a vertex incident to the edge which has the critical point;*
(6) *exchanging the relative values of two vertices* v *and* v' *such that there is no pair of edges* e *and* e' *in the subdivided graph* $\cup_i \Gamma_i$ *such that* e *is incident to* v*,* e' *is incident to* v' *and* e *and* e' *share a common vertex;*

Examples of these are shown in Figure 9.

Proof. This follows essentially immediately from elementary Cerf theory for 1-manifolds with boundary, though there are several comments we need to make. Given two Morse functions f_0, f_1 on Γ (the unsubdivided graph) we consider a path of smooth functions f_t between f_0 and f_1. Generically such a function f_t will be Morse except at finitely many t. Similarly generically f_t will also only have finitely many t where the vertices of Γ and the critical points of f_t do not have distinct values.

Suppose that $\mathcal{C}_0$ and $\mathcal{C}_1$ are Cerf decompositions obtained by taking a collection of values

$$0 = s_0^i < \cdots < s_{n_i}^i = 1$$

such that $f_i^{-1}([s_j^i, s_{j+1}^i])$ contains at most one critical point of f_i or one vertex of Γ. By applying moves of type (1), we can assume that there is exactly one critical point of f_i or vertex of Γ in each $f_i^{-1}([s_j^i, s_{j+1}^i])$.

Now if there are no points $t \in [a, b]$ where f_t has not Morse and all of the critical values and vertex values of f are distinct, then let $c_1^t, \ldots, c_n^t$ denote the critical values and vertex values of f, ordered so that $c_1^t < \cdots c_n^t$. Pick values $s_j^t \in [0, 1]$ extending s_j^i for $i \in \{0, 1\}$ such that

$$0 = s_0^t < c_1^t < s_1^t < \cdots < c_n^t < s_n^t = 1.$$

The sets $\Gamma_{i,t} = f_t^{-1}([s_i^t, s_{i+1}^t)$ yield the underlying graph for a Cerf decomposition of $\mathcal{G}$. If $\mathcal{C}_t$ denotes the Cerf decomposition of $\mathcal{G}$ arising from the subsets $\Gamma_{i,t}$, then it is easy to see that the $\mathcal{C}_a$ and $\mathcal{C}_b$ are related by a sequence of moves of type (1).

Suppose a birth-death singularity of f_t occurs at a point $t = c \in (0, 1)$. Then if there are no other singularities of f_t in $[c - \epsilon, c + \epsilon]$, and $\mathcal{C}_{c-\epsilon}, \mathcal{C}_{c+\epsilon}$ are Cerf decompositions as above for $f_{c-\epsilon}$ and $f_{c+\epsilon}$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ are related by a composition of moves of type (1) and moves of type (2).

Similarly if a critical point is destroyed or created along on edge at a vertex incident to that edge at a time $c \in (0, 1)$, and there are no other singularities in $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ are related by a move of type (3) and moves of type (1).

Suppose two critical points exchange relative value at a time $c \in (0, 1)$, and there are no other singularities in $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ are related by a move of type (4) and moves of type (1) as long as the edges containing the critical two points do not contain a common vertex. Observe that if the edges with a critical point contain a common vertex (of the subdivided graph), then at $t = c$ the derivative at all of the points between the two critical points must be zero, which is a non-generic singularity. Hence generically the two critical points switching relative values corresponds to a composition of moves (4) and (1).

If a critical point and a vertex exchange relative value at a point $c \in (0, 1)$ and there are no other singularities in $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ relative by a move of type (5) and moves of type (1) so long as there is no edge which is both incident to both the vertex v and one of the two vertices on the edge with a critical point. Note that once again, if there were such an edge, then at the derivative of f_c along this edge must vanish, which is a nongeneric singularity. Hence generically we can assume that there is no such edge (of the subdivided graph).

Finally if two vertices v and v' (of the nonsubdivided graph Γ) exchange relative values and there are no other singularities in the interval $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}, \mathcal{C}_{c+\epsilon}$ are related by moves (1) and (6), so long as

there are not edges e and e' in the subdivided graph such that e is incident to v, e' is incident to v', and e and e' contain a common vertex. As in the previous cases, if there were two such edges e and e', then f_c would have vanishing derivative on all of $e \cup e'$, which is a nongeneric singularity.

We note that these are the only singularities that occur generically for a 1-parameter family of functions f_t defined on a graph Γ.

We finally comment that for a Morse function to yield a Cerf decomposition as we've described, we must have $\mathbf{w}_0 \cup f_t^{-1}(s)$ nonempty for any s. If $\mathbf{w}_0 \neq \varnothing$, this is trivially satisfied. In the case that $\mathbf{w}_0 = \varnothing$, given two Morse functions on a flow graph which both have nonempty level sets, it is possible to connect them via a generic path of functions which have nonempty level sets. This can be achieved by pulling a small portion of a terminal edge down through the level sets and fixing the value on this portion through the homotopy of Morse functions, and then moving it back into its final position.

□

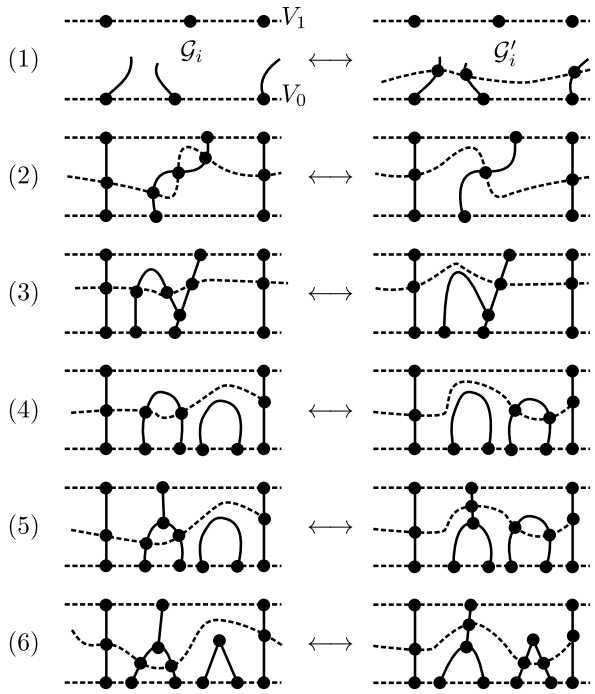

FIGURE 9. Six moves which can be used to move between any two Cerf decompositions of a strong ribbon flow graph. All of the dashed lines indicate level sets of Morse functions, and the space between them indicates an elementary strong ribbon flow graph. The dashed lines may be somewhat misleading, because strictly speaking the Morse function is only defined on the graph.

We now construct the graph action map. Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is an embedded strong ribbon flow graph in Y^3 and that $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. Suppose that $V \subseteq V(\Gamma)$ is a distinguished set of vertices with $V_0 \subseteq V$.

We first define the graph action map for an elementary strong ribbon flow graph. Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is an elementary strong ribbon flow graph of type (1), $(3i)$, or $(3ii)$. Then we define

$$A^V_{\mathcal{G},\mathfrak{s}} = \prod_{v \in V(\Gamma)\setminus V_1} S^-_{v,\mathfrak{s}} \circ \prod_{e \in E(\Gamma)} A_{e,\mathfrak{s}} \circ \prod_{v \in V(\Gamma)\setminus V} S^+_{v,\mathfrak{s}}.$$

Notice that this formula is well defined because no two edges share a common vertex, so by Lemma 5.9 we know that the order of the various A_e is not important. Similarly the S^+_v and S^-_v can be commuted amongst each other.

Now suppose that $\mathcal{G}$ is an elementary strong ribbon flow graph of type 2. Let v denote the unique vertex which is not in V_0 or V_1. Let $e_1, \dots, e_n$ be the edges adjacent to v with the ordering from Γ. We define the graph action map by the formula

$$A^V_{\mathcal{G},\mathfrak{s}} = \left(\prod_{v \in V(\Gamma) \setminus V_1} S^-_{v,\mathfrak{s}} \right) \circ \left(\prod_{e \in E(\Gamma) \setminus \{e_1,\dots,e_n\}} A_{e,\mathfrak{s}} \right) \circ (A_{e_n,\mathfrak{s}} \cdots A_{e_1,\mathfrak{s}}) \circ \left(\prod_{v \notin V} S^+_{v,\mathfrak{s}} \right).$$

Now given an arbitrary strong ribbon flow graph $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ on Y^3, and a collection $V \subseteq V(\Gamma)$ such that $V_0 \subseteq V$, we pick a Cerf decomposition $(\mathcal{G}_1, \dots, \mathcal{G}_n)$ of $\mathcal{G}$, and define

$$A^V_{\mathcal{G},\mathfrak{s}} : HF^\circ(Y, V \cup \mathbf{w}_0, \mathfrak{s}, \sigma|_{V \cup \mathbf{w}_0}) \to HF^\circ(Y, V_1 \cup \mathbf{w}_0, \mathfrak{s}, \sigma|_{V_1 \cup \mathbf{w}_0})$$

as the composition

$$A^V_{\mathcal{G},\mathfrak{s}} = \prod_{i=1}^{n} A^{V_{\geq i}}_{\mathcal{G}_i,\mathfrak{s}},$$

where

$$V_{\geq i} = V \cap \bigcup_{j \geq i} V(\mathcal{G}_j)$$

Theorem 7.6. *If $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is a strong ribbon flow graph with V, then the map $A_{\mathcal{G}}$ is independent of the Cerf decomposition $(\mathcal{G}_1, \dots, \mathcal{G}_n)$.*

We need to show that the composition $A_{\mathcal{G}}$ is unchanged by moves (1) through (6) in Lemma 7.5. Before we prove the above theorem, we will first state a model computation that we will use repeatedly, but will prove later:

Lemma 7.7. *Suppose that e is a path from v to v'. Then there exists a diagram and an almost complex structure such that, viewing CF° as freely stabilized at v, we have that*

$$A_e = \begin{pmatrix} 0 & U_v \\ 1 & 0 \end{pmatrix},$$

and furthermore this diagram and almost complex structure can be used to compute S^+_v and S^-_v. In particular

$$A_e^2 = U_v$$

and

$$S^-_v A_e S^+_v = 1.$$

The proof will be given later, as Lemma 14.6. The latter two relations follow easily from the first claim. Note that $A_e^2 \simeq U_v$ was previously proven in Lemma 5.11, by counting the ends of moduli spaces of Maslov index 2-disks.

The following lemma states essentially that the graph action map is invariant under subdivision of an edge.

Lemma 7.8. *Suppose e_1 is a path from w to w' and e_2 is a path from w' to w''. Let $A_{e_2 * e_1}$ denote the relative homology map on a diagram which does not have w' as a basepoint. Then we have*

$$A_{e_2 * e_1} = S^-_{w'} A_{e_2} A_{e_1} S^+_{w'}.$$

Proof. By Lemma 7.7 we have $S^-_{w'} A_{e_1} S^+_{w'} = 1$. Hence we have that

$$\begin{aligned} A_{e_2 * e_1} &= A_{e_2 * e_1}(S^-_{w'} A_{e_1} S^+_{w'}) \\ &= S^-_{w'} A_{e_2 * e_1} A_{e_1} S^+_{w'} \\ &= S^-_{w'} (A_{e_2} + A_{e_1}) A_{e_1} S^+_{w'} \\ &= S^-_{w'} A_{e_2} A_{e_1} S^+_{w'} + S^-_{w'} U_{w'} S^+_{w'} \\ &= S^-_{w'} A_{e_2} A_{e_1} S^+_{w'}, \end{aligned}$$

as we wanted. □

We need another lemma, about subdividing an edge into thirds. From here on, let $S^+_{w_1 w_2 \cdots w_n}$ abbreviate $S^+_{w_1} \circ \cdots \circ S^+_{w_n}$ and similarly for the destabilization maps.

Lemma 7.9. *Suppose e_1, e_2 and e_3 are edges from w_0 to w_1, from w_1 to w_2, and from w_2 to w_3, respectively, as is depicted in Figure 10. Suppose also that $\sigma : \{1,2,3\} \to \{1,2,3\}$ is a bijection, then*

$$S^-_{w_1 w_2} A_{e_{\sigma(3)}} A_{e_{\sigma(2)}} A_{e_{\sigma(1)}} S^+_{w_1 w_2} = S^-_{w_1 w_2} A_{e_3} A_{e_2} A_{e_1} S^+_{w_1 w_2}.$$

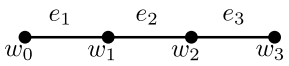

FIGURE 10. Vertices and edges in Lemma 7.9.

Proof. We will show that if e_i, e_j and e_k are the edges, then

$$S^-_{w_1 w_2} A_{e_i} A_{e_j} A_{e_k} S^+_{w_1 w_2} = S^-_{w_1 w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1 w_2}.$$

The same argument will apply to show that the expression is invariant under switching the order of the left two relative homology maps as well.

There are two cases: if e_j and e_k share a vertex, or if they do not. If e_i and e_j do not share a vertex, then $A_{e_j} A_{e_k} = A_{e_k} A_{e_j}$ so the above equality trivially holds.

Now suppose that e_j and e_k share the vertex $w_n \in \{w_1, w_2\}$. Let w_m denote the other vertex in $\{w_1, w_2\}$. We simply compute that

$$\begin{aligned} S^-_{w_1 w_2} A_{e_i} A_{e_j} A_{e_k} S^+_{w_1 w_2} &= S^-_{w_1 w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1 w_2} + S^-_{w_1 w_2} A_{e_i} U_{w_n} S^+_{w_1 w_2} \\ &= S^-_{w_1 w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1 w_2} + S^-_{w_m} U_{w_n} S^-_{w_n} S^+_{w_n} A_{e_i} S^+_{w_m} \\ &= S^-_{w_1 w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1 w_2}. \end{aligned}$$

□

We now have the necessary tools to prove the invariance of the graph action map:

Proof of Theorem 7.6. **Move (1):** In the case that the elementary flow graph $\mathcal{G}$ is of type (1), or (2). This follows from Lemma 7.8. In the case that the graph is of type $(3i)$ or $(3ii)$, invariance follows from Lemmas 7.8 and 7.9.

Move (2): Suppose that v_1, v_2, v_3, v_4, and v_5 are vertices, and e_1, e_2, e_3 and e_4 are edges, as in Figure 11.

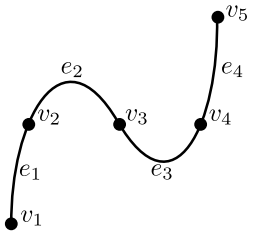

FIGURE 11. Vertices and edges used in the proof of invariance under move (2).

Write e for the concatenation of e_1, e_2, e_3 and e_4. To prove invariance under move (2), it is sufficient to show that

$$S^-_{v_1} A_e = (S^-_{v_2 v_3 v_4} A_{e_4} A_{e_2})(S^-_{v_1} A_{e_3} A_{e_1} S^+_{v_2 v_3 v_4}).$$

To this end, we compute as follows:

$$\begin{aligned} S^-_{v_1} A_e &= S^-_{v_1 v_2 v_3 v_4} A_{e_4} A_{e_3} A_{e_2} A_{e_1} S^+_{v_2 v_3 v_4} && \text{(Lemma 7.8)} \\ &= S^-_{v_1 v_2 v_3 v_4} A_{e_4} A_{e_2} A_{e_3} A_{e_1} S^+_{v_2 v_3 v_4} && \text{(Lemma 7.9)} \\ &= S^-_{v_2 v_3 v_4} A_{e_4} A_{e_2} S^-_{v_1} A_{e_3} A_{e_1} S^+_{v_2 v_3 v_4} && \text{(Lemma 6.6),} \end{aligned}$$

as we wanted.

Move (3): We consider move (3) shown in Figure 9. There is another move which we would need to check obtained by flipping this picture upside down, but we leave checking that move to the reader since the proof is identical to what we will present. Let v, v', v'' and $e_2, \dots, e_n, e, e'$, and e'' be as in Figure 12. We let e_n be the concatenation of e, e' and e''.

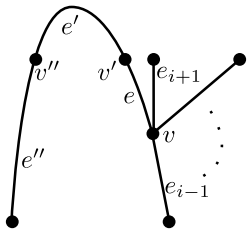

FIGURE 12. Vertices and edges used in the proof of invariance under move (3). Here $e_i = e * e' * e''$.

To show invariance under move (3), it is sufficient to show that

$$A_{e_n} \cdots A_{e_1} S_v^+ = (S^-_{v'v''} A_{e'})(A_{e''} A_{e_n} \cdots A_{e_{i+1}} A_e A_{e_{i-1}} \cdots A_{e_1} S^+_{v'v''}).$$

To this end, we compute:

$$\begin{aligned}
A_{e_n} \cdots A_{e_i} \cdots A_{e_1} &= A_{e_n} \cdots A_{e_{i+1}} (S^-_{v'} A_{e'*e''} A_e S^+_{v'}) A_{e_{i-1}} \cdots A_{e_1} && \text{(Lemma 7.8)} \\
&= A_{e_n} \cdots A_{e_{i+1}} (S^-_{v'} (S^-_{v''} A_{e'} A_{e''} S^+_{v''}) A_e S^+_{v'}) A_{e_{i-1}} \cdots A_{e_1} && \text{(Lemma 7.8)} \\
&= S^-_{v''v'} A_{e_n} \cdots A_{e_{i+1}} A_{e''} A_{e'} A_e A_{e_{i-1}} \cdots A_{e_1} S^+_{v'v''} && \text{(Lemma 6.6)} \\
&= (S^-_{v''v'} A_{e'})(A_{e''} A_{e_n} \cdots A_{e_{i+1}} A_e A_{e_{i-1}} \cdots A_{e_1} S^+_{v'v''}) && \text{(Lemma 5.9)}
\end{aligned}$$

as we wanted.

Moves (4)-(6): Invariance under moves (4)-(6) follows similarly from Lemmas 7.8, 7.9 and 5.9. □

Having established the independence of the graph action map from the choice of Cerf decomposition, we now prove two lemmas about the graph action map which will be useful later when we construct and analyze the graph cobordism maps. Firstly, the graph action map is invariant under subdivision, in the following sense:

Lemma 7.10. *Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is a strong ribbon flow graph in Y^3 and $\mathcal{G}'$ is a subdivision of $\mathcal{G}$ formed by adding a single vertex v to the interior of an edge in Γ. Then if V is a collection of vertices containing V_0 with $v \notin V$, then*

$$A^V_{\mathcal{G}} = A^{V \cup \{v\}}_{\mathcal{G}'} \circ S^+_v,$$

for either absolute ordering of the two edges adjacent to v.

Proof. Pick a Cerf decomposition of $\mathcal{G}$ such that the vertex v appears in an elementary flow graph $\mathcal{G}_v$ of type (1) (a trivial flow graph). Let e_v denote the edge of $\mathcal{G}_v$ containing v. Let e_1 and e_2 denote the two components of $e_v \setminus \{v\}$, with either ordering. By Corollary 10.3 we know that

$$A_{e_v} = S^-_v A_{e_1} A_{e_2} S^+_v.$$

The left hand side A_{e_v} is a factor in $A^V_{\mathcal{G}}$. If we insert the right hand expression into the formula for $A^V_{\mathcal{G}}$ and then commute S^+_v with everything to the right, we get exactly $A^{V \cup \{v\}}_{\mathcal{G}'} \circ S^+_v$. □

We now consider the effect of adding a "trivial strand" to a strong ribbon flow graph:

Lemma 7.11. *Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is a strong ribbon flow graph in Y and let Γ_0 denote a graph obtained by adding a new vertex $v \in Y \setminus \Gamma$ to Γ and also adding an edge connecting v to a point on the interior of an edge of Γ. Let $\mathcal{G}_0 = (\Gamma_0, V_0, V_1, \sigma, \mathbf{w}_0)$, $\mathcal{G}' = (\Gamma_0, V_0 \cup \{v\}, V_1, \sigma, \mathbf{w}_0)$ and $\mathcal{G}'' = (\Gamma_0, V_0, V_1 \cup \{v\}, \sigma, \mathbf{w}_0)$. Then*

$$A_{\mathcal{G}_0} = A_{\mathcal{G}},$$
$$S_v^- \circ A_{\mathcal{G}'} = A_{\mathcal{G}},$$

and

$$A_{\mathcal{G}''} \circ S_v^+ = A_{\mathcal{G}}.$$

This holds for any ordering of the vertices adjacent to v.

Proof. We will first show that $A_{\mathcal{G}_0} = A_{\mathcal{G}}$ with the absolute ordering of e_1, e_2, λ. By taking a Cerf decomposition of $\mathcal{G}$ which we have subdivided enough, one easily computes that it is sufficient to show that

$$S_v^- S_{v_2}^- A_\lambda A_{e_2} A_{e_1} S_{v_2}^+ S_v^+ = S_{v_2}^- A_{e_2} A_{e_1} S_{v_2}^+,$$

where $v, v_1, v_2, \lambda, e_1$ and e_2 are as in Figure 13.

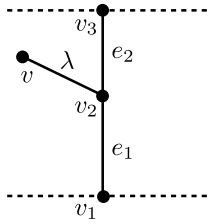

FIGURE 13. Edges and vertices in the proof of Lemma 7.11. The dashed lines indicate the incoming and outgoing ends of the flow graph.

By Lemma 7.7, we have that $S_v^- A_\lambda S_v^+ = 1$, and hence we have

$$\begin{aligned} S_{v_2}^- A_{e_2} A_{e_1} S_{v_2}^+ &= S_{v_2}^- (S_v^- A_\lambda S_v^+) A_{e_2} A_{e_1} S_{v_2}^+ \\ &= S_v^- S_{v_2}^- A_\lambda A_{e_2} A_{e_1} S_{v_2}^+ S_v^+, \end{aligned}$$

as we wanted. We conclude that $A_{\mathcal{G}} = A_{\mathcal{G}_0}$. To see that the statement holds for any other absolute ordering of the edges e_1, e_2, and λ, we observe that in the above computation

$$S_{v_2}^- A_{e_2} A_{e_1} S_{v_2}^+ = S_{v_2}^- A_{e_1} A_{e_2} S_{v_2}^+,$$

allowing us to reorder e_1 and e_2. Also the expression $S_v^- A_\lambda S_v^+$ can be inserted anywhere in $S_{v_2} A_{e_1} A_{e_2} S_{v_2}^+$, allowing us to reorder λ and either of the edges e_1 and e_2.

We now consider $A_{\mathcal{G}'}$ and $A_{\mathcal{G}''}$. We note that we've established that $A_{\mathcal{G}} = A_{\mathcal{G}_0}$ for any Cerf decomposition of $\mathcal{G}_0$. In particular, a Cerf decomposition of $\mathcal{G}'$ naturally induces a Cerf decomposition of $\mathcal{G}_0$. The only difference in the two formulas defining $A_{\mathcal{G}'}$ and $A_{\mathcal{G}_0}$ is a factor of S_v^- which is present in $A_{\mathcal{G}_0}$ but not in $A_{\mathcal{G}'}$ and hence

$$S_v^- \circ A_{\mathcal{G}'} = A_{\mathcal{G}_0},$$

establishing the desired equality.

Similarly to show that $A_{\mathcal{G}''} \circ S_v^+ = A_{\mathcal{G}_0}$, we pick a Cerf decomposition of $\mathcal{G}''$, which naturally induces one for $\mathcal{G}_0$. The only difference in the two formulas defining these operations is a factor of S_v^+ to the left of all other terms in $A_{\mathcal{G}''}$ but not in $A_{\mathcal{G}_0}$, establishing the desired equality.

□

We finally show that the graph action map is invariant of cyclic reordering of the edges incident to a vertex $v \in V(\Gamma)$ as long as $v \notin V$, the initial set of vertices of the graph action map.

Lemma 7.12. *Suppose that $e_1, \ldots, e_n$ are edges which each share a vertex v (and v has valence 1 in each edge e_i). Then*

$$S_v^- A_{e_n} \cdots A_{e_2} A_{e_1} S_v^+ = S_v^- A_{e_{n-1}} \cdots A_{e_1} A_{e_n} S_v^+.$$

Proof. We first prove the cases $n = 1, 2$ and $n = 3$ directly. The cases $n > 3$ will follow by an induction argument. First note that $n = 1$ is vacuous. The $n = 2$ is relatively easy. We just compute directly that

$$S_v^- A_{e_2} A_{e_1} S_v^+ = S_v^- U_v S_v^+ + S_v^- A_{e_1} A_{e_2} S_v^+ = S_v^- A_{e_1} A_{e_2} S_v^+,$$

by Lemma 5.9. The $n = 3$ is slightly more involved, though still relatively easy. We simply compute that

$$\begin{aligned} S_v^- A_{e_2} A_{e_1} A_{e_0} S_v^+ &= S_v^- A_{e_1} A_{e_2} A_{e_0} S_v^+ + U_v S_v^- A_{e_0} S_v^+ \\ &= S_v^- A_{e_1} A_{e_0} A_{e_2} S_v^+ + U_v S_v^- (A_{e_0} + A_{e_2}) S_v^+ \\ &= S_v^- A_{e_1} A_{e_0} A_{e_2} S_v^+ + U_v A_{e_1 * e_0} S_v^- S_v^+ \\ &= S_v^- A_{e_1} A_{e_0} A_{e_2} S_v^+. \end{aligned}$$

We now prove the case of $n > 3$. This is more involved. First pick an extra vertex w which is not an endpoint of any of the e_i and a path e from v to w. We first claim that

$$A_{e_2} A_{e_1} = S_w^- A_{e_2 * e} A_{e_1 * e} A_e S_w^+.$$

To see this, note that $A_{e_2 * e} = A_{e_2} + A_e$. Hence we just compute that

$$\begin{aligned} S_w^- A_{e_2 * e} A_{e_1 * e} A_e S_w^+ &= S_w^- (A_{e_2} + A_e)(A_{e_1} + A_e) A_e S_w^+ \\ &= S_w^- A_{e_2} A_{e_1} A_e S_w^+ + S_w^- (A_{e_2} A_e^2 + A_e A_{e_1} A_e + A_e^3) S_w^+ \\ &= A_{e_2} A_{e_1} S_w^- A_e S_w^+ + S_w^- (A_{e_2} U_w + A_{e_1} U_w + U_w A_e + U_w A_e) S_w^+ \\ &= A_{e_2} A_{e_1} + U_w S_w^- A_{e_2 * e_1} S_w^+ \\ &= A_{e_1} A_{e_1} + U_w A_{e_2 * e_1} S_w^- S_w^+ \\ &= A_{e_1} A_{e_1}, \end{aligned}$$

as we wanted.

Let $\widetilde{e}_1 = e_1 * e$ and let $\widetilde{e}_2 = e_2 * e$. We compute as follows:

$$\begin{aligned} S_v^- A_{e_n} \cdots A_{e_2} A_{e_1} S_v^+ &= S_v^- A_{e_n} \cdots A_{e_3} (S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_e S_w^+) S_v^+ \\ &= S_v^- S_w^- (A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_{e_n} \cdots A_{e_3} A_e) S_w^+ S_v^+ \\ &= S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_n} \cdots A_{e_3} A_e S_v^+) S_w^+. \end{aligned} \tag{7}$$

By induction, Equation (7) is equal to

$$\begin{aligned} S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_n} \cdots A_{e_3} A_e S_v^+) S_w^+ &= S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_e A_{e_n} \cdots A_{e_3} S_v^+) S_w^+ \\ &= S_v^- S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_e S_w^+ A_{e_n} \cdots A_{e_3} S_v^+ \\ &= S_v^- A_{e_2} A_{e_1} A_{e_n} \cdots A_{e_3} S_v^+. \end{aligned}$$

This implies that the expression $S_v^- A_{e_n} \cdots A_{e_1} S_v^+$ is invariant under permuting the e_i cyclically by two elements at a time.

Similarly, by induction we know that Equation (7) is also equal to

$$\begin{aligned} S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_n} \cdots A_{e_3} A_e S_v^+) S_w^+ &= S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_3} A_e A_{e_n} \cdots A_{e_4} S_v^+) S_w^+ \\ &= S_v^- A_{e_3} (S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_e S_w^+) A_{e_3} A_{e_n} \cdots A_{e_4} S_v^+ \\ &= S_v^- A_{e_3} A_{e_2} A_{e_1} A_{e_n} \cdots A_{e_4} S_v^+, \end{aligned}$$

which shows that the expression $S_v^- A_{e_n} \cdots A_{e_1} S_v^+$ is invariant under permuting the e_i cyclically three elements at a time. Since it is also invariant under permuting the e_i cyclically two elements at a time, we know that it is invariant under permuting them cyclically one at a time, so the main claim now follows. □

Theorem B now follows from Theorem 7.6 Lemma 7.12.

8. 1- and 3-handles

In this section we construct and prove invariance of 1- and 3-handle maps. Our construction of the maps is very similar to the ones in [OS06] and [Juh15], though it has some notable differences. Unlike the construction in [OS06], our construction works for 1-handles with feet in different components, allowing us to consider cobordisms between manifolds with disconnected ends. Our construction is somewhat more flexible than the one in [Juh15], since we don't require the feet of our 1-handles to be in a neighborhood of the basepoints, a trick that eliminates most of the need for gluing lemmas for the hat flavor, but which isn't useful for the $-, \infty$ and $+$ flavors.

8.1. **Preliminaries on 1- and 3-handles.** Suppose that $p_1, p_2 \in Y$ are two points. By choosing an orientation on $\{p_1, p_2\}$, we can view $\{p_1, p_2\}$ as a framed zero sphere $\mathbb{S}_0 \subseteq Y$. Surgery on $\mathbb{S}_0$ corresponds to removing two disjoint balls containing p_1 and p_2, then gluing a copy of $S^2 \times I$ connecting the two resulting boundary components. We call the resulting 3-manifold $Y(\mathbb{S}_0)$, though it depends on the choice of balls containing the points p_1 and p_2 (though of course any two choices of balls yield manifolds which are diffeomorphic with diffeomorphism which is unique up to isotopy). In terms of cobordisms, we can construct a cobordism $W(\mathbb{S}_0)$ with boundary $-Y \sqcup Y(\mathbb{S}_0)$ by attaching a 1-handle to $Y \times I$ along $\mathbb{S}_0 \times \{1\}$.

Dually if Y has an embedded 2-sphere $\mathbb{S}_2$ with trivial normal bundle, then after removing a regular neighborhood of $\mathbb{S}_2$, we get a three manifold with two boundary components, both of which are 2-spheres. By filling in the boundary components with balls, we get a closed three manifold $Y(\mathbb{S}_2)$. In terms of cobordisms, we can form a cobordism $W(\mathbb{S}_2)$ by attaching a 3-handle to $Y \times I$ along $\mathbb{S}_2 \times \{1\}$.

Definition 8.1. *Given an embedded 0-sphere $\mathbb{S}_0 = \{p_1, p_2\} \subseteq Y \setminus \mathbf{w}$ and a choice of open balls $\mathcal{B}_1, \mathcal{B}_2 \subseteq Y \setminus \mathbf{w}$ such that $p_i \in \mathcal{B}_i$ and $\mathcal{B}_1 \cap \mathcal{B}_2 = \varnothing$, we say that a Heegaard splitting $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is **subordinate** to $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ if $p_1, p_2 \in \Sigma$, and $\mathcal{B}_i \cap \Sigma$ is a disk which is disjoint from $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w}$.*

Given a Spin^c structure $\mathfrak{s}$ on Y, and a framed sphere $\mathbb{S}_0$, there is a unique Spin^c structure $\overline{\mathfrak{s}}$ on $W(\mathbb{S}_0)$ extending $\mathfrak{s}$. Let $\mathfrak{s}'$ denote the restriction of $\overline{\mathfrak{s}}$ to $Y(\mathbb{S}_0)$. The map $\mathfrak{s} \mapsto \mathfrak{s}'$ is an injection. Dually, given a framed 2-sphere $\mathbb{S}_2 \subseteq Y$ and a Spin^c structure $\mathfrak{s}' \in \text{Spin}^c(Y)$, if $\mathfrak{s}'$ extends over $W(\mathbb{S}_2)$, then it does so uniquely. We let $\overline{\mathfrak{s}}$ denote the extension over $W(\mathbb{S}_2)$ and we let $\mathfrak{s}$ denote the restriction of $\overline{\mathfrak{s}}$ to $Y(\mathbb{S}_2)$.

We will define maps

$$F^\circ_{Y,\mathbb{S}_0,\overline{\mathfrak{s}}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}) \to HF^\circ(Y(\mathbb{S}_0), \mathbf{w}, \mathfrak{s}')$$

and

$$F^\circ_{Y,\mathbb{S}_2,\overline{\mathfrak{s}}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}') \to HF^\circ(Y(\mathbb{S}_2), \mathbf{w}, \mathfrak{s}).$$

If we let $\mathbb{S}_0 = \{p_1, p_2\} \subseteq \mathbb{S}_2$ where each p_i is in one of the balls that we used to fill in the boundary components of $Y \setminus N(\mathbb{S}_2)$, we clearly have $Y(\mathbb{S}_2)(\mathbb{S}_0) = Y$, so instead of defining maps from $HF^\circ(Y, \mathbf{w}, \mathfrak{s})$ to $HF^\circ(Y(\mathbb{S}_2), \mathbf{w}, \mathfrak{s}')$ for embedded 2-spheres, it is equivalent to define maps $HF^\circ(Y(\mathbb{S}_0), \mathbf{w}, \mathfrak{s}') \to HF^\circ(Y, \mathbf{w}, \mathfrak{s})$ for embedded 0-spheres $\mathbb{S}_0$, which is the strategy we will take.

Now given a Heegaard diagram $\mathcal{H}$ which is subordinate to $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$, and some additional data $\mathcal{J}$ (gluing data) which we will define in the next subsection, we will define maps

$$F^\circ_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\overline{\mathfrak{s}}} : CF^\circ_{J_s}(\mathcal{H}, \mathfrak{s}) \to CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_\mathcal{J}, \mathfrak{s}')$$

and

$$G^\circ_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\overline{\mathfrak{s}}} : CF_{\mathcal{J}(T)}(\mathcal{H}_\mathcal{J}, \mathfrak{s}') \to CF^\circ_{J_s}(\mathcal{H}, \mathfrak{s}).$$

If $\mathcal{H}$ is the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$, then $\mathcal{H}_\mathcal{J} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \mathbf{w})$ is a diagram formed by removing disks from Σ centered at p_1 and p_2, then connecting the corresponding boundary components in $Y(\mathbb{S}_0)$ by an annulus. We define $\overline{\boldsymbol{\alpha}} = \boldsymbol{\alpha} \cup \{\alpha_0\}$ and $\overline{\boldsymbol{\beta}} = \boldsymbol{\beta} \cup \{\beta_0\}$, where α_0 and β_0 are homotopically nontrivial, simple closed curves in the annulus which are isotopies of one another and satisfy $|\alpha_0 \cap \beta_0| = 2$.

Note that on the level of complexes, we have the following:

$$CF^\circ(\mathcal{H}_\mathcal{J}, \mathfrak{s}') = CF^\circ(\mathcal{H}, \mathfrak{s}) \otimes_{\mathbb{Z}_2} \mathbb{Z}_2^2 = CF^\circ(\mathcal{H}, \mathfrak{s}) \oplus CF^\circ(\mathcal{H}, \mathfrak{s}).$$

We will define the 1-handle cobordism maps by

$$F^\circ_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\overline{\mathfrak{s}}}(\mathbf{x}) = \mathbf{x} \times \theta^+$$

where θ^+ is the top degree intersection point of $\alpha_0 \cap \beta_0$. We will define the 3-handle map by

$$G^\circ_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\bar{\mathfrak{s}}}(\mathbf{x}\times\theta^+) = 0, \qquad \text{and} \qquad G^\circ_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\bar{\mathfrak{s}}}(\mathbf{x}\times\theta^-) = \mathbf{x},$$

for appropriate choice of almost complex structure.

Lemma 8.2. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for $(Y, \mathbf{w})$ which is subordinate to the zero sphere $\mathbb{S}_0 = \{p, p'\} \subseteq \Sigma$. Let $\overline{\mathcal{H}}$ be the diagram described above for $Y(\mathbb{S}_0)$. Suppose $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ and that $\mathfrak{s}' \in \mathrm{Spin}^c(Y(\mathbb{S}_0))$ is the Spin^c structure corresponding to $\mathfrak{s}$ under the correspondence in the previous paragraphs. Then $\mathcal{H}$ is strongly $\mathfrak{s}$-admissible iff $\overline{\mathcal{H}}$ is strongly $\bar{\mathfrak{s}}$-admissible.*

Proof. Suppose $\mathcal{H}$ is strongly $\mathfrak{s}$-admissible. Suppose $\mathcal{D} = \mathcal{D}_1 \# \mathcal{D}_2$ is a periodic domain on $\overline{\mathcal{H}}$ with $\langle c_1(\bar{\mathfrak{s}}), H(\mathcal{D})\rangle = 2N$, where $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ and $\mathcal{D}_2$ is a periodic domain on our standard diagram for $(S^2, \alpha_0, \beta_0, w, w')$. Using Lemmas 4.7 and 6.3, an easy computation shows that

$$\langle c_1(\bar{\mathfrak{s}}), H(\mathcal{D})\rangle = \langle c_1(\mathfrak{s}), H(\mathcal{D}_1)\rangle,$$

so $\mathcal{D}_1$ has a multiplicity greater than N, and hence so does $\mathcal{D}$. It follows that $\overline{\mathcal{H}}$ is strongly $\bar{\mathfrak{s}}$-admissible.

Conversely if $\overline{\mathcal{H}}$ is strongly $\bar{\mathfrak{s}}$-admissible and $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ with $\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 2N$. Let A_w and $A_{w'}$ be the components of $S^2 \setminus \alpha_0$ containing w and w' respectively. Let $\mathcal{D}_2$ be the domain on $(S^2, \alpha_0, \beta_0, w, w')$ defined by

$$\mathcal{D}_2 = n_p(\mathcal{D}_1)A_w + n_{p'}(\mathcal{D}_1)A_{w'}.$$

Since $\overline{\mathcal{H}}$ is strongly $\bar{\mathfrak{s}}$-abmissible, we know that $\mathcal{D}_1 \# \mathcal{D}_2$ has coefficients strictly bigger than N, and hence so does $\mathcal{D}_1$, so $\mathcal{H}$ is strongly $\mathfrak{s}$-admissible. □

8.2. **Gluing data for 1- and 3-handle maps.** Very similarly to the free stabilization maps, we now define "gluing data" for the 1- and 3-handle maps. The only real difference is that we are gluing at two points instead of one. In the following definition, let $(S^2, \alpha_0, \beta_0, w, w')$ be the diagram for (S^3, w, w') shown in Figure 5.

Definition 8.3. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(Y, \mathbf{w})$ which is subordinate to a framed sphere $\mathbb{S}_0 = \{p_1, p_2\} \subseteq Y \setminus \{w\}$. We call a collection*

$$\mathcal{J} = (\mathcal{A}, J_s, J'_s, B_1, B'_1, B_2, B'_2, r_1, r'_1, r_2, r'_2, p'_1, p'_2, \iota, \phi_1, \phi_2)$$

gluing data for $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ where

1. *$\mathcal{A}$ is a choice of annulus in $S^2 \times I \subseteq Y(\mathbb{S}_0)$ extending the circles $\Sigma \cap \partial\mathcal{B}_i$ such that $\mathcal{A} \cap (S^2 \times \{t\})$ is a circle for all t;*
2. *$\mathbb{S}_0 = \{p_1, p_2\} \subseteq Y$ is a 0-sphere and $\mathcal{B}_i$ is an open ball in Y containing p, such that $Y(\mathbb{S}_0)$ is formed by gluing a copy of $S^2 \times I$ to $Y \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$.*
3. *B_i is a closed ball on Σ which doesn't intersect any $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w}$ and which contains the point $p_i \in \Sigma$ and which contains $\mathcal{B}_i \cap \Sigma$, and such that $B_1 \cap B_2 = \varnothing$;*
4. *the points p'_i are on S^2 and p'_1 is in the same region as w' and p'_2 is in the same region as w.*
5. *B'_i is a closed ball on S^2 containing p'_i which doesn't intersect $\alpha_0 \cup \beta_0 \cup \{w, w'\}$, and $B'_1 \cap B'_2 = \varnothing$;*
6. *J_s is an almost complex structure on $\Sigma \times I \times \mathbb{R}$ which is split on B_1 and B_2;*
7. *J'_s is an almost complex structure on $S^2 \times I \times \mathbb{R}$ which is split on B'_1 and B'_2;*
8. *r_1, r_2 are real numbers such that $0 < r_1, r_2 < 1$ and such that $\mathcal{B}_i \cap \Sigma$ is contained in $\mathrm{int}(r_i \cdot B_i)$;*
9. *r'_1, r'_2 are two real numbers such that $0 < r'_1, r'_2 < 1$;*
10. *$\iota : \overline{S^2 \setminus (r'_1 \cdot B'_1 \cup r'_2 \cdot B'_2)} \hookrightarrow \mathcal{A}$ is an embedding;*
11. *if $\widetilde{A}_i$ denotes the annulus inside of $(\Sigma \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)) \cup \widetilde{A}$ with boundary equal to $\partial B_i \cup \partial B'_i$, then*

$$\phi_i : \widetilde{A}_i \to S^1 \times [-a_i, 1 + b_i]$$

is a diffeomorphism which restricts to a conformal map of $B_i \setminus (r_i \cdot B_i)$ to $S^1 \times [-a_i, 0]$ and a conformal map of $B'_i \setminus (r'_i \cdot B'_i)$ to $S^1 \times [1, 1 + b_i]$.

As in the case of the free stabilization maps, gluing data $\mathcal{J}$ determines a diagram $\mathcal{H}_\mathcal{J}$ as well as an almost complex structure $\mathcal{J}(T_1, T_2)$ for $T_i > 1$ which have necks length T_1 and T_2.

8.3. **Counting holomorphic disks after surgery on $\mathbb{S}_0$.** In this section, we prove most of the holomorphic curve components of the proof of naturality of the 1- and 3-handle maps. Formally the 1- and 3-handle maps are nearly identical to the free stabilization maps. This is essentially because the free stabilization maps are secretly a 0- or 4-handle map composed with a 1- or 3-handle map, a fact which we will uncover later when we consider subdivision invariance of the graph cobordism maps.

The strategy will be very similar to our analysis of the free stabilization map, the only real difference that we are counting disks which satisfy a matching condition at two points, instead of one. Let $(S^2, \alpha_0, \beta_0, w, w')$ be the diagram in Figure 5. Let p_1 be a point in the same region as w and let p_2 be a point in the same region as w'. Given a homotopy class of disks $\phi \in \pi_2(a, b)$ for $a, b \in \{\theta^+, \theta^-\}$ and a pair of divisors $\mathbf{d}_1 \in \mathrm{Sym}^{k_1}(\mathbb{D})$ and $\mathbf{d}_2 \in \mathrm{Sym}^{k_2}(\mathbb{D})$ with $\mu(\phi) = 2k_1 + 2k_2$, we define

$$\mathcal{M}(\phi, \mathbf{d}_1, \mathbf{d}_2) = \{u \in \mathcal{M}(\phi) : \rho^{p_1}(u) = \mathbf{d}_1, \rho^{p_2}(u) = \mathbf{d}_2\}.$$

The following lemma is very similar to [OS08, Lem. 6.4.]:

Lemma 8.4. *Let $(S^2, \alpha_0, \beta_0, w, w')$ and p_1 and p_2 be points in the same regions as w and w' respectively. Let $\mathbf{d}_1 \in \mathrm{Sym}^{k_1}(\mathbb{D})$ and $\mathbf{d}_2 \in \mathrm{Sym}^{k_2}(\mathbb{D})$ be two divisors. Write $k = k_1 + k_2$. Then*

$$\sum_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi) = 2k}} \#\mathcal{M}(\phi, \mathbf{d}_1, \mathbf{d}_2) \equiv 1 (\mathrm{mod}\, 2).$$

Proof. We will first show that the quantity

$$M(a, \mathbf{d}_1, \mathbf{d}_2) \stackrel{\mathrm{def}}{=} \sum_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi) = 2k}} \#\mathcal{M}(\phi, \mathbf{d}_1, \mathbf{d}_2)$$

is independent modulo 2 of $\mathbf{d}_1$ and $\mathbf{d}_2$. This follows from the following Gromov compactness argument. Suppose that $(\mathbf{d}_1^0, \mathbf{d}_2^0)$ and $(\mathbf{d}_1^1, \mathbf{d}_2^1)$ are two pairs of divisors and that $\mathbf{d}_i^t$ is a path from $\mathbf{d}_i^0$ to $\mathbf{d}_i^1$. Consider the 1-dimensional moduli space

$$\mathcal{M}(a, \mathbf{d}_1, \mathbf{d}_2) = \bigcup_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi) = 2k}} \mathcal{M}(\phi, \mathbf{d}_1^t, \mathbf{d}_2^t).$$

There are four kinds of ends of this space. Since $\mathcal{M}(a, \mathbf{d}_1, \mathbf{d}_2)$ can be compactified to a compact 1-manifold, the total number of ends is zero. The first two types correspond to setting $t = 0$ or $t = 1$, i.e. we have one end of the form

$$\bigcup_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi) = 2k}} \mathcal{M}(\phi, \mathbf{d}_1^0, \mathbf{d}_2^0)$$

and another of the form

$$\bigcup_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi) = 2k}} \mathcal{M}(\phi, \mathbf{d}_1^1, \mathbf{d}_2^1).$$

The other two ends correspond to a degenerations of a collections of homologous holomorphic strips into a broken holomorphic strip, i.e. collections of holomorphic strips, boundary degenerations, and closed surfaces. Since one of the components must satisfy a matching condition with a pair of divisors $(\mathbf{d}_1^t, \mathbf{d}_2^t)$ for some t, using the Maslov index calculation in Lemma 6.3, we can constrain the types of degenerations which can occur, and we are left with the following two types of ends:

$$\bigcup_{\substack{\phi_1 \in \pi_2(a,b) \\ \phi_2 \in \pi_2(b,a) \\ t \in I \\ \mu(\phi_1) = 1 \\ n_p(\phi_1) = n_{p'}(\phi_1) = 0 \\ \mu(\phi_2) = 2k-1}} \widehat{\mathcal{M}}(\phi_1) \times M(\phi_2, \mathbf{d}_1^t, \mathbf{d}_2^t)$$

and

$$\bigcup_{\substack{\phi_1\in\pi_2(a,b)\\ \phi_2\in\pi_2(b,a)\\ t\in I\\ \mu(\phi_1)=2k-1\\ n_{p_1}(\phi_2)=n_{p_2}(\phi_2)=0\\ \mu(\phi_2)=1}} M(\phi_1,\mathbf{d}_1^t,\mathbf{d}_2^t)\times\widehat{\mathcal{M}}(\phi_2).$$

Each of the above two spaces have an even number of ends because the differential on $\widehat{CF}(S^2,\alpha_0,\beta_0,w,w')$ vanishes. Hence modulo 2, the total number of ends is

$$M(a,\mathbf{d}_1^0,\mathbf{d}_2^0)-M(a,\mathbf{d}_1^1,\mathbf{d}_2^2)$$

and hence the number $M(a,\mathbf{d}_1,\mathbf{d}_2)$ is independent of the divisors $\mathbf{d}_i$.

To show that $M(a\mathbf{d}_1,\mathbf{d}_2)\equiv 1(\mathrm{mod}\,2)$, pick divisors $\mathbf{d}_1^T$ and $\mathbf{d}_2^T$ such that $\mathbf{d}_1^T\cup\mathbf{d}_2^T$ is a collection with vertical spacing at least T between each pair of points and as $T\to\infty$ the points in $\mathbf{d}_1^T\cup\mathbf{d}_2^T$ all approach $\{0\}\times\mathbb{R}\subseteq\mathbb{D}$. As we let $T\to\infty$, one argues as in [OS08, Lem. 6.4] that for sufficiently large T, any such strip is equal to k_1+k_2 boundary degenerations glued to the constant strip, and hence $M(a\mathbf{d}_1^T,\mathbf{d}_2^T)$ is 1 for sufficiently large T. □

Lemma 8.5. *Suppose that $\mathbb{S}_0=\{p_1,p_2\}\subseteq Y$ is an embedded 0-sphere and that $\mathcal{H}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w})$ is a strongly $\mathfrak{s}$-admissible diagram for $(Y,\mathbf{w})$ which is subordinate to $\mathbb{S}_0$. Suppose that $\mathcal{J}$ is gluing data with almost complex structure J_s on $\mathcal{H}$. For sufficiently large T and T' we have an identification of differentials*

$$\partial_{\mathcal{J}(T,T')}=\begin{pmatrix}\partial_{J_s} & *\\ 0 & \partial_{J_s}\end{pmatrix}.$$

Proof. The proof is very similar to the proof of Lemma 6.5. Write

$$\partial_{J_s(T)}=\begin{pmatrix}A & B\\ C & D\end{pmatrix}.$$

We first compute A and D. Suppose that $\phi=\phi_1\#\phi_2$ is a Maslov index one homology disk and $\phi\in\pi_2(\mathbf{x}\times\theta^+,\mathbf{y}\times\theta^+)$ or $\phi\in\pi_2(\mathbf{x}\times\theta^-,\mathbf{y}\times\theta^-)$, which admits holomorphic representatives $u_{T,T'}$ for T and T' which are both arbirarily large. Then by extracting a convergent subsequence, we know that ϕ_1 and ϕ_2 each have representatives U_1 and U_2 as broken holomorphic strips. Now using the Maslov index computation in Lemma 6.3 as well as the combinatorial formula for the Maslov index, we see that

$$\mu(\phi)=\mu(\phi_1)+\mu(\phi_2)-2m_1-2m_2=\mu(\phi_1).$$

We conclude that $\mu(\phi_1)=1$. In particular, the broken holomorphic strip U_1 can have at most one nonconstant component, u_1. There must be a component u_2 of U_2 satisfying the matching condition

$$\rho^{p_1}(u_1)=\rho^{p_1'}(u_2)\qquad\text{and}\qquad\rho^{p_2}(u_1)=\rho^{p_2'}(u_2).$$

We now claim that the remaining components of U_2 must be constant. To see this, let u_2 represent the homology class ϕ_2'. Since $\phi_1\#\phi_2'$ has a holomorphic representative $u_1\#u_2$ given by gluing, we know that $\mu(\phi_1\#\phi_2')\geq 1$. From the matching condition, we know that $n_p(\phi_2-\phi_2')=n_{p'}(\phi_2-\phi_2')=0$. Hence an easy Maslov index computation shows that

$$1=\mu(\phi)=\mu(\phi_1\#\phi_2')+\mu(0\#(\phi_2-\phi_2'))=\mu(\phi_1\#\phi_2')+\mu(\phi_2-\phi_2'),$$

from which it follows that $\mu(\phi_2-\phi_2')=0$, and hence $\phi_2-\phi_2'$ has no nonconstant components since it admits a representative as a broken holomorphic curve. Hence we conclude that if both T and T' are sufficiently large, then any holomorphic strip u representing $\phi_1\#\phi_2$ is of the form $u_1\#u_2$ where u_1 and u_2 are preglued holomorphic curves.

Conversely, given a Maslov index 1 homology class ϕ_1 on $\mathcal{H}$ which admits a holomorphic representative u_1, by Lemma 8.4, we know that among homology classes $\phi_2\in\pi_2(\theta^+,\theta^+)$, there is a one (modulo 2) holomorphic representative u_2 such that

$$\rho^{p_1}(u_1)=\rho^{p_1'}(u_2)\qquad\text{and}\qquad\rho^{p_2}(u_1)=\rho^{p_2'}(u_2).$$

The same statement holds for $\phi_2\in\pi_2(\theta^-,\theta^-)$. Combining this with the previous paragraph yields an identification $A=D=\partial_{J_s}$.

We now wish to show that $B = 0$. Suppose that $\phi = \phi_1 \# \phi_2 \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^-)$ is a Maslov index one homology disk. Note that by Lemma 6.3, we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 - 2m_2 = \mu(\phi_1) + 1.$$

Hence $\mu(\phi_1) = 0$. If ϕ has representatives for T and T' both arbitrarily large, then ϕ_1 and ϕ_2 have broken holomorphic representatives U_1 and U_2. Since $\mu(\phi_1) = 0$, we conclude that U_1 consists only of constant components. Hence

$$n_{p_1}(\phi_1) = n_{p_2}(\phi_1) = n_{p_1'}(\phi_2) = n_{p_2'}(\phi_2) = 0.$$

By direct examination, ϕ_2 is limited to one of two homology disks. Each has a unique holomorphic representative, which cancel modulo 2. Hence $C = 0$. □

Similar to the free stabilization maps, we also need to show that the relative homology maps commute with the 1- and 3-handle maps.

Lemma 8.6. *Suppose that λ is a path in Y and $\mathbb{S}$ is a framed, embedded 0- or 2-sphere. Suppose that $\overline{\mathfrak{s}}$ is a* Spin^c *structure on $W(Y, \mathbb{S})$. Assume that λ is an embedded path in Y which doesn't intersect $\mathbb{S}$. Let $\overline{\lambda}$ be the induced path in $Y(\mathbb{S})$. The following relation holds:*

$$A_{\overline{\lambda}} \circ F^\circ_{Y,\mathbb{S},\overline{\mathfrak{s}}} = F^\circ_{Y,\mathbb{S},\overline{\mathfrak{s}}} \circ A_\lambda.$$

Proof. Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ be a diagram for $(Y, \mathbf{w})$ which is subordinate to the 1-handle data $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ and suppose that $\mathcal{J}$ is gluing data for $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$. Let $\mathcal{H}_\mathcal{J}$ denote the resulting diagram for $Y(\mathbb{S}_0)$. Suppose that λ is a curve between two basepoints in $\mathbf{w}$ which doesn't intersect the points $\mathbb{S}_0$, and let $\overline{\lambda}$ denote the induced curve in $Y(\mathbb{S}_0)$. It is sufficient to show that

$$A_{\overline{\lambda}} = \begin{pmatrix} A_\lambda & * \\ 0 & A_\lambda \end{pmatrix},$$

from which the main claim follows. This follows similar to the proof of Lemma 8.5, but with an additional factor of $a(\lambda, \phi)$ on the counts of holomorphic disks. □

Lemma 8.7. *Suppose $\mathbb{S} \subseteq Y$ is a framed 0- or 2-sphere and $\mathbf{w}$ is a collection of basepoints for Y with $\mathbf{w} \cap \mathbb{S} = \varnothing$. If $w \in Y \setminus (\mathbf{w} \cup \mathbb{S})$, then*

$$S^\pm_w \circ F^\circ_{Y,\mathbb{S},\mathfrak{s}} = F^\circ_{Y,\mathbb{S},\mathfrak{s}} \circ S^\pm_w$$

for a $\mathfrak{s} \in$ Spin$^c(W(Y, \mathbb{S}))$.

Proof. This is follows immediately from the formulas for $S^\pm_w$ and $F^\circ_{Y,\mathbb{S},\mathfrak{s}}$, so long as an almost complex structure can be picked which puts the maps in their standard forms. It is not hard to argue that this can be done (cf. the subclaim at the end of Theorem 10.2). □

8.4. **Counting holomorphic triangles after surgery on $\mathbb{S}_0$.** As we did in Section 6.4 for the free stabilization maps, we would like to prove an analogous stabilization result for triangle maps after adding a 1-handle. Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple which is strongly admissible for a Spin^c structure $\mathfrak{s}$ and which is subordinate to the 1-handle data $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ with $\mathbb{S}_0 = \{p_1, p_2\}$. Let $\mathcal{J}$ be gluing data for the triple $\mathcal{T}$. We form the 1-handle stabilized Heegaard triple $\overline{\mathcal{T}} = \{\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}}, \mathbf{w})$ where $\overline{\Sigma}$ is the surface obtained by removing $\Sigma \cap (\mathcal{B}_1 \cup \mathcal{B}_2)$, and connecting the resulting boundary components with an annulus as specified by $\mathcal{J}$. Also we define

$$\overline{\boldsymbol{\alpha}} = \boldsymbol{\alpha} \cup \{\alpha_0\}$$

and similarly for $\overline{\boldsymbol{\beta}}$ and $\overline{\boldsymbol{\gamma}}$, where α_0, β_0 and γ_0 are homotopicically nontrivial simple closed curves in the annulus which are all pairwise isotopic to one another, and such that

$$|\alpha_0 \cap \beta_0| = |\beta_0 \cap \gamma_0| = |\alpha_0 \cap \beta_0| = 2,$$

analogous to Figure 7.

In light of Lemma 4.11, we have

$$\mathfrak{s}_\mathbf{w}(\psi_1 \# \psi_2) = \mathfrak{s}_\mathbf{w}(\psi_1' \# \psi_2') \qquad \text{iff} \qquad \mathfrak{s}_\mathbf{w}(\psi_1) = \mathfrak{s}_\mathbf{w}(\psi_1'),$$

if ψ_1 is a triangle on $\mathbb{T}$ and ψ_2 is a triangle on $(S^2, \alpha_0, \beta_0, \gamma_0, w, w')$. In particular a Spin^c structure $\mathfrak{s} \in$ Spin$^c(X_{\alpha\beta\gamma})$ uniquely determines a $\overline{\mathfrak{s}} \in$ Spin$^c(X_{\overline{\alpha}\overline{\beta}\overline{\gamma}})$ such that $\mathfrak{s}_\mathbf{w}(\psi_1 \# \psi_2) = \overline{\mathfrak{s}}$ iff $\mathfrak{s}_\mathbf{w}(\psi_1) = \mathfrak{s}$.

We have the following stabilization theorem, analogous to Theorem 6.17.

Theorem 8.8. *For sufficiently large T, the stabilized triangle maps $\overline{\mathcal{T}}$ take the form*

$$F_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^+)) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix},$$

where the matrix on the right is an expansion of the map on the left in the variable $\cdot$, which can be either x^+ or x^-. The maps $F^\circ_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}$ are computed with respect to the almost complex structure $\mathcal{J}(T)$, and $F^\circ_{\mathcal{T},\mathfrak{s}}$ is computed with the almost complex structure J_s on Σ.

Similarly, for sufficiently large T, we have

$$F_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}((\mathbf{x}\times x^+)\otimes(\mathbf{y}\times\cdot)) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix}.$$

As with the analysis of the free stabilization maps, to prove Theorem 8.8, we need to perform a count of holomorphic triangles which satisfy a certain matching condition. Suppose that $\psi\in\pi_2(x,y,z)$ is a homology triangle on $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$, with $x\in\alpha_0\cap\beta_0, y\in\beta_0\cap\gamma_0$ and $z\in\alpha_0\cap\gamma_0$. Suppose $\mathbf{d}_1\in\mathrm{Sym}^{k_1}(\mathbb{D})$ and $\mathbf{d}_2\in\mathrm{Sym}^{k_2}(\mathbb{D})$. We then define

$$\mathcal{M}(\psi,\mathbf{d}_1,\mathbf{d}_2) = \{u\in\mathcal{M}(\psi) : \rho^{p_1'}(u)=\mathbf{d}_1, \rho^{p_2'}(u)=\mathbf{d}_2\},$$

and

$$M_{(x,y,z)}(\mathbf{d}_1,\mathbf{d}_2) = \sum_{\substack{\psi\in\pi_2(x,y,z)\\ \mu(\psi)=2k_1+2k_2}} \#\mathcal{M}(\psi,\mathbf{d}_1,\mathbf{d}_2).$$

Lemma 8.9. *If (x,y,z) is one of the triples $(x^+,y^+,z^+),(x^-,y^+,z^-)$ or (x^+,y^-,z^-), then for generic $\mathbf{d}_1$ and $\mathbf{d}_2$ we have*

$$M_{(x,y,z)}(\mathbf{d}_1,\mathbf{d}_2)\equiv 1(\mathrm{mod}\,2).$$

Proof. The proof is identical to the proof of Lemma 6.19. □

Proof of Theorem 8.8. With Lemma 8.9, the proof is now follows essentially identically to the proof of 6.17. □

8.5. **Dependence on gluing data.** Similar to the free stabiliation maps, we have the following:

Lemma 8.10. *Suppose that $\mathcal{H}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w})$ is a strongly $\mathfrak{s}$-admissible diagram for $(Y,\mathbf{w})$ and let $(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2)$ be 1-handle data in Y where $\mathbb{S}_0=\{p_1,p_2\}$ is a framed 0-sphere in Y. Suppose that $\mathcal{J}$ is gluing data for $(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2)$. Write J_s for the almost complex structure on $\Sigma\times I\times\mathbb{R}$ specified by $\mathcal{J}$. Then there is an N such that if $T,T',T'',T'''>N$, and all almost complex structures achieve transversality, then the map*

$$\Phi_{\mathcal{J}(T,T')\to\mathcal{J}(T'',T''')} : CF^\circ_{\mathcal{J}(T,T')}(\mathcal{H}_{\mathcal{J}},\mathfrak{s}') \to CF^\circ_{\mathcal{J}(T'',T''')}(\mathcal{H}_{\mathcal{J}},\mathfrak{s}')$$

is homotopy equivalent to

$$\Phi_{\mathcal{J}(T,T')\to\mathcal{J}(T'',T''')} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. The proof is identical to the proof of Lemma 6.8, though we use Lemma 8.4 to find preglued disks. □

Definition 8.11. *We say that N is **sufficiently large** for gluing data $\mathcal{J}$, if $\Phi_{\mathcal{J}(T,T')\to\mathcal{J}(T'',T''')}$ is of the form in the previous lemma for all $T,T',T'',T'''\geq N$. We say T is **large enough to compute** the 1- or 3-handle maps if $T>N$ for some N which is sufficiently large.*

Lemma 8.12. *Suppose that $\mathcal{J}_1$ and $\mathcal{J}_2$ are two choices of gluing data for $(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2)$ and that $\mathcal{H}$ is strongly $\mathfrak{s}$-admissible and subordinate to $(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2)$. Then if T_i and T_i' are large enough to the compute the 1- and 3-handle maps for $\mathcal{J}_i$ using the above definition, then*

$$\Phi_{\mathcal{J}_1(T_1,T_1')\to\mathcal{J}_2(T_2,T_2')} \simeq \begin{pmatrix} \Phi_{J_{s,1}\to J_{s,2}} & * \\ 0 & \Phi_{J_{s,1}\to J_{s,2}} \end{pmatrix}$$

where $J_{s,i}$ is the almost complex structure on $\Sigma\times I\times\mathbb{R}$ associated to $\mathcal{J}_i$.

Proof. The effect of changing all of the pieces of gluing data except for the annulus $\mathcal{A}$ filling in $\Sigma \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$ follows nearly identically to the analogous result for the free stabilization maps. One simply verifies that the proofs of Lemmas 6.10, 6.11, 6.12, 6.13, 6.14, 6.15, and 6.16 adapt essentially without change from the free stabiliation setting to the 1- and 3-handle setting. Then one verifies that the above moves are sufficient to move between $\mathcal{J}_1$ to $\mathcal{J}_2$ as long as we use the same annulus, $\mathcal{A}$, for both $\mathcal{J}_1$ and $\mathcal{J}_2$ is the same.

The annulus $\mathcal{A}$ is specified up to isotopy by the following argument: inside of $Y(\mathbb{S}_0)$ there is a subset identified with $S^2 \times I$. By assumption $\mathcal{A} \cap (S^2 \times \{t\})$ is a single circle. Such an embedding of $S^1 \times I$ into $S^2 \times I$ always extends to an embedding of $\mathbb{D}^2 \times I$ into $S^2 \times I$, and any two embeddings are isotopic to each other via an isotopy fixing $S^2 \times \partial I$ since S^2 is simply connected.

Hence given two annuli $\mathcal{A}$ and $\widetilde{\mathcal{A}}$ which fill in $\Sigma \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$ in $Y(\mathbb{S}_0)$, there is an isotopy ψ_t of $Y(\mathbb{S}_0)$ which is fixed outside of $Y \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$ which maps $\mathcal{A}$ to $\widetilde{\mathcal{A}}$. If $\mathcal{J}$ is gluing data, then we can move to gluing data $\widetilde{\mathcal{J}} = (\psi_1)_* \mathcal{J}$, which has annulus $\widetilde{\mathcal{A}}$ and the change of almost complex structures map $\Phi_{\mathcal{J}(T,T') \to \widetilde{\mathcal{J}}(T,T')}$ is tautologically the isotopy map $(\psi_1)_*$, which is the identity matrix once we identity the two complexes. □

8.6. **Invariance of 1- and 3-handle maps.** In this section we prove invariance of the 1- and 3-handle maps.

Theorem 8.13. *Suppose that $(W, \mathbf{w})$ is a based 3-manifold with 1-handle data $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$. Suppose that $\overline{\mathfrak{s}} \in \mathrm{Spin}^c(W(\mathbb{S}_0))$, $\mathfrak{s} = \overline{\mathfrak{s}}|_Y$, and suppose that $\mathcal{H}_1$ and $\mathcal{H}_2$ are two strongly $\mathfrak{s}$-admissible diagrams for Y. Suppose that $\mathcal{J}_1$ and $\mathcal{J}_2$ are choices of gluing data for $\mathcal{H}_1$ and $\mathcal{H}_2$ respectively. Then the following diagram and its analog for the 3-handle maps commutes up to chain homotopy:*

$$\begin{array}{ccc} CF^\circ_{J_{s,1}}(\mathcal{H}_1) & \xrightarrow{F_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J}_1,\overline{\mathfrak{s}}}} & CF^\circ_{\mathcal{J}_1(T_1)}((\mathcal{H}_1)_{\mathcal{J}_1}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}_1 \to \mathcal{H}_2}} & & \downarrow{\scriptstyle \Phi_{(\mathcal{H}_1)_{\mathcal{J}_1} \to (\mathcal{H}_2)_{\mathcal{J}_2}}} \\ CF^\circ_{J_{s,2}}(\mathcal{H}_2) & \xrightarrow{F_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J}_2,\overline{\mathfrak{s}}}} & CF^\circ_{\mathcal{J}_2(T_2)}((\mathcal{H}_2)_{\mathcal{J}_2}). \end{array}$$

Proof. The proof of the analogous statement for the free stabilization maps (Theorem 6.22) holds in this context essentially without change. Invariance from gluing data follows from Lemma 8.12. Invariance from $\boldsymbol{\alpha}$- and $\boldsymbol{\beta}$-handleslides and isotopies follows from Theorem 8.8. Invariance from isotopies of Y which fix the basepoints $\mathbf{w}$ follows identically to Theorem 6.22. □

9. 2-handle maps

In this section we construct and analyze the 2-handle maps associated to an embedded framed link $\mathbb{L} \subseteq Y$. These maps were essentially defined in their current form by Ozsváth and Szabó in [OS06], though not in the cylindrical setting, and not in the case of multiple basepoints. Our contribution to these maps is essentially limited to gluing lemmas which we use to show that the 2-handle maps cancel the 1- and 3-handle maps in the appropriate sense, as well as showing that the 2-handle maps commute with the free stabilization, relative homology, and 1- and 3-handle maps under appropriate circumstances. In addition to the original construction in [OS06, Sec. 4], we will rely heavily on work from [Juh15, Sec. 8], which builds the theory for 3-manifolds with multiple basepoints.

9.1. **Maps associated to a framed link.**

Definition 9.1. *A **framed link** $\mathbb{L}$ in a 3-manifold Y is a collection of n pairwise disjoint embedded circles $K_1, \dots, K_n \subseteq Y$ together with a **framing**, i.e. a choice of homology classes $\ell_i \in H_1(\partial N(K_i))$ satisfying $m_i \cdot \ell_i = 1$ where m_i is the meridian of K_i.*

Definition 9.2. *A **bouquet** $B(\mathbb{L})$ for a link $\mathbb{L}$ is an embedded 1-complex in Y which is the union of $\mathbb{L}$ with a collection of arcs $a_1, \dots, a_n$ where each a_i is an arc from a component K_i of $\mathbb{L}$ to a basepoint in $\mathbf{w}$. The arcs a_i do not all need to connect to the same basepoint. We define F_i to be the punctured torus $\partial N(K_i) \setminus N(a_i)$.*

Definition 9.3. *We say a Heegaard triple*

$$\mathbb{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}) = (\Sigma, \{\alpha_1, \dots, \alpha_{g+\ell-1}\}, \{\beta_1, \dots, \beta_{g+\ell-1}\}, \{\gamma_1, \dots, \gamma_{g+\ell-1}\}, \mathbf{w}),$$

*is **subordinate to the bouquet** $B(\mathbb{L})$ if*

(1) $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ *is an embedded Heegaard diagram for* $(Y, \mathbf{w})$*;*
(2) $(\Sigma, \{\alpha_1, \dots, \alpha_{g+\ell-1}\}, \{\beta_{n+1}, \dots, \beta_{g+\ell-1}\}, \mathbf{w})$ *is a diagram for* $(Y \setminus N(B(\mathbb{L}))$*. By this we mean the following: write* U_α *and* U_β *for the closures of the components of* $Y \setminus \Sigma$ *and let* $\mathbf{w}^\alpha$ *denote small pushoffs of the basepoints* $\mathbf{w}$ *into* U_α*. Then the* $\alpha_1, \dots, \alpha_{g+\ell+1}$ *curves bound compressing disks and after compressing* Σ *along these disks in* U_α *we get a surface which is isotopic to the spheres* $\partial N(\mathbf{w}^\alpha))$*. Let* $(B(\mathbb{L}) \cup \mathbf{w})^\beta$ *denote a pushoff of* $B(\mathbb{L}) \cup \mathbf{w}$ *into* U_β*. The curves* $\beta_{n+1}, \dots, \beta_{g+\ell-1}$ *bound compressing disks in* $U_\beta \setminus N(B(\mathbb{L}))$*, and after compressing* Σ *along these disks, we get a surface which is isotopic to* $\partial N((B(\mathbb{L}) \cup \mathbf{w})^\beta)$*;*
(3) the curves $\gamma_{n+1}, \dots, \gamma_{g+\ell-1}$ *are small isotopies of the curves* $\beta_{n+1}, \dots, \beta_{g+\ell-1}$*;*
(4) after compressing Σ *along the curves* $\beta_{n+1}, \dots, \beta_{g+\ell-1}$*, the induced curves* β_i *and* γ_i *(for* $i = 1, \dots, n$*) lie in the punctured torus* F_i*;*
(5) for $i = 1, \dots, n$ *the curve* β_i *represents a meridian of* K_i *and* $\beta_i \cap \gamma_j = \varnothing$ *unless* $i = j$*, in which case* β_i *and* γ_i *meet at a single transverse intersection point;*
(6) for $i = 1, \dots, n$ *the homology class of* γ_i *corresponds to the framing* ℓ_i *under an identification of*

$$H_1(\partial N(\cup_i K_i)) \cong H_1(\partial N(B(\mathbb{L}))).$$

Definition 9.4. *If* $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ *is a Heegaard triple subordinate to* $B(\mathbb{L}) \subseteq Y$*, and* $\mathcal{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$ *is another Heegaard triple subordinate to* $B(\mathbb{L})$*, we say that* $\mathcal{T}'$ *is the* ***stabilization*** *of* $\mathcal{T}$ *if* Σ' *is the internal connected sum of* Σ *and a genus one surface* $E \subseteq Y \setminus N(B(\mathbb{L}))$ *with* $E \cap \Sigma = \varnothing$ *and* $\boldsymbol{\tau}' = \boldsymbol{\tau} \cup \{\tau_0\}$ *for* $\tau \in \{\alpha, \beta, \gamma\}$ *where* α_0, β_0 *and* γ_0 *are curves on* E *such that* $|\alpha_0 \cap \beta_0| = 1$ *and* β_0 *and* γ_0 *are small isotopies of each other with* $|\beta_0 \cap \gamma_0| = 2$*.*

The following is analogous to [OS06, Prop. 4.3]:

Lemma 9.5. *Suppose that* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ *is a Heegaard triple subordinate to the bouquet* $B(\mathbb{L})$ *in* Y*. Then* $X_{\alpha\beta\gamma}$ *is a cobordism from* Y *to*

$$Y(\mathbb{L}) \sqcup (\sqcup_i \#^{k_i}(S^1 \times S^2))$$

where k_i *is the number of* $\boldsymbol{\alpha}$ *curves in the* i^{th} *component of* Σ*. The monodiagram* $(\Sigma, \boldsymbol{\beta})$ *defines a cobordism* X_β *from* $U_\beta \cup_\Sigma (-U_\beta)$ *to* $\varnothing$*. Let* $W_\beta(\mathbb{L})$ *denote the cobordism obtained from* W_β *by attaching 4-dimensional 2-handles to* $\mathbb{L}$ *along* U_β*. The manifold* $W_\beta(\mathbb{L})$ *is a cobordism from* $Y_{\beta\gamma}$ *to* $\varnothing$*. We can glue* $X_{\alpha\beta\gamma}$ *and* $W_\beta(\mathbb{L})$ *along* $Y_{\beta\gamma}$ *to obtain a cobordism equivalent to* $W(\mathbb{L})$*, and this equivalence is well-defined up to isotopy fixing* $\mathbf{w} \times I \subseteq W(\mathbb{L})$*.*

Proof. This follows from [Juh15, Prop. 6.6]. □

Lemma 9.6. *Suppose* $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ *is subordinate to* $B(\mathbb{L})$*, and* $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$*. Under the inclusion* $Y_{\beta\gamma} \subseteq X_{\alpha\beta\gamma} \subseteq W(\mathbb{L})$*, the* Spin^c *structure* $\mathfrak{s}|_{Y_{\beta\gamma}}$ *is the torsion* Spin^c *structure.*

Proof. This follows as in [OS06, Prop. 4.3]. The cobordism $W(\mathbb{L})$ is diffeomorphic to $X_{\boldsymbol{\alpha\beta\gamma}}$ with 3- and 4-handles attached along $Y_{\beta\gamma}$. The only Spin^c structure extending over these 3- and 4-handles is the torsion Spin^c structure. □

We can now define the two handle maps. Given a framed link $\mathbb{L} \subseteq Y$, a bouquet $B(\mathbb{L})$, an $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$, and a Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ which is subordinate to $B(\mathbb{L})$ and which is strongly $\mathfrak{s}|_{X_{\alpha\beta\gamma}}$-admissible, we can define

$$F^\circ_{Y,\mathbb{L},\mathcal{T},\mathfrak{s}} = HF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s}|_Y) \to HF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{L})})$$

by the formula

$$F^\circ_{Y,\mathbb{L},\mathcal{T},\mathfrak{s}}(\mathbf{x}) = F^\circ_{\mathcal{T},\mathfrak{s}|_{X_{\alpha\beta\gamma}}}(\mathbf{x} \otimes \Theta^+),$$

where $\Theta^+ \in CF^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathfrak{s}_0)$ is the top degree generator. We now show invariance of this map from the Heegaard triple $\mathcal{T}$, and then from the bouquet $B(\mathbb{L})$. We will address independence of this map from the bouquet and the intermediate Heegaard diagram, but first we must prove several versions of $(1,2)$-stabilization invariance of the triangle maps.

9.2. **(1,2)-stabilization and triangle maps.** We now prove $(1,2)$-stabilization invariance of the triangle maps. In principle this is "standard" since it is needed to prove many foundational results, for example the invariance of the 3-manifold invariants, but to the authors knowledge a proof of this fact in the cylindrical context has not appeared in the literature. The following argument proceeds along the same lines as Theorem 6.17, first performing a count of triangles on a simple diagram with a certain matching condition, and then performing a stretching argument to show that all triangles on the stabilized diagram limit to preglued triangles as we let the neck length approach ∞. Let $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ be the diagram shown in Figure 14. There is a new feature of this argument, which is not present in the neck stretching arguments we performed for free stabilization and for the 1- and 3-handle maps. To avoid closed holomorphic surfaces from appearing, which do not necessarily achieve transversality, we need to consider somewhat more harshly perturbed almost complex structures, similar to those satisfying $(J5')$ in [Lip06].

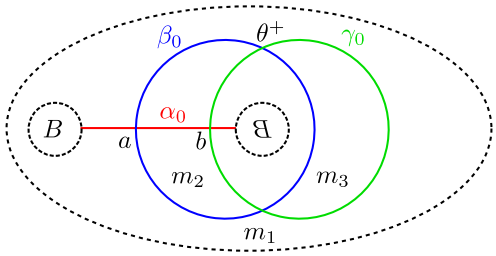

FIGURE 14. The Heegaard Triple $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ used to perform $(1,2)$-stabilization of triangle maps.

Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is strongly admissible and let $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ be the triple shown in Figure 14. Given an $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma}})$, then in light of Lemma 4.11, $\mathfrak{s}$ determines a unique $\overline{\mathfrak{s}}$ in $\mathrm{Spin}^c(X_{\overline{\boldsymbol{\alpha}}\overline{\boldsymbol{\beta}}\overline{\boldsymbol{\gamma}}})$ for the $(1,2)$-stabilized triple $\overline{\mathcal{T}} = (\Sigma \# \mathbb{T}^2, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}}, \mathbf{w})$, where $\overline{\boldsymbol{\tau}} = \boldsymbol{\tau} \cup \{\tau_0\}$ for $\tau \in \{\alpha, \beta, \gamma\}$. Using this correspondence between Spin^c structures, we have the following:

Theorem 9.7. *Suppose $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a strongly $\mathfrak{s}$-admissible with a distinguished point $q \in \Sigma \setminus (\mathbf{w} \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma})$ and $\overline{\mathcal{T}}$ is the Heegaard triple*

$$\overline{\mathcal{T}} = (\Sigma \# \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w}),$$

where the connected sum is taken at $p \in \mathbb{T}^2$ and $q \in \Sigma$. Then $\overline{\mathcal{T}}$ is strongly $\overline{\mathfrak{s}}$-admissible and if θ^+ is the top generator of $\beta_0 \cap \gamma_0$, then for sufficiently stretched almost complex structure (which is suitably generic in a sense described below) there is an identification of the $\mathbf{z} \times b$ component of

$$F^{\circ}_{\overline{\mathcal{T}}, \overline{\mathfrak{s}}}((\mathbf{x} \times a) \otimes (\mathbf{y} \times \theta^+))$$

and the $\mathbf{z}$ component of

$$F^{\circ}_{\mathcal{T}, \mathfrak{s}}(\mathbf{x} \otimes \mathbf{y})$$

where $\{a\} = \alpha_0 \cap \beta_0$ and $\{b\} = \alpha_0 \cap \gamma_0$.

Suppose $\overline{\mathcal{T}}'$ instead denotes the triple defined similarly to $\overline{\mathcal{T}}$, but instead with α_0 and β_0 isotopic and with $|\alpha_0 \cap \beta_0| = 2$ and $|\alpha_0 \cap \gamma_0| = |\beta_0 \cap \gamma_0| = 1$. Letting θ^+ denote the top degree generator of $\alpha_0 \cap \beta_0$ and $\{a\} = \beta_0 \cap \gamma_0$ and $\{b\} = \alpha_0 \cap \gamma_0$, then for sufficiently stretched and generic almost complex structure there is an identification of the $\mathbf{z} \times b$ component of

$$F^{\circ}_{\overline{\mathcal{T}}', \mathfrak{s}}((\mathbf{x} \times \theta^+) \otimes (y \times a))$$

and the $\mathbf{z}$ component of

$$F^{\circ}_{\mathcal{T}, \mathfrak{s}}(\mathbf{x} \times \mathbf{y}).$$

We will only prove the above claim for $\overline{\mathcal{T}}$, since the proof for $\overline{\mathcal{T}}'$ is nearly identical. We also remark that admissibility is easy to check, so we leave it to the reader (cf. Lemma 8.2). There is a somewhat subtle issue in that to make the correct statements, we must work with almost complex structures for which closed surfaces mapped into the interior of Δ are non generic. In [Lip06], Lipshitz considers almost complex

structures on $\Sigma \times I \times \mathbb{R}$ satisfying five axioms, labelled $(J1), \ldots, (J5)$. For proving stabilization invariance of the chain complexes, he considers an alternate fifth axiom $(J5')$ which is slightly more generic, and for which closed holomorphic curves mapped into the interior of Δ are non generic. In constructing the holomorphic triangle maps, he considers almost complex structures on $\Sigma \times \Delta$ satisfying axioms labelled $(J'1) - (J'4)$. Following Lipshitz's approach to stabilization invariance of the chain complexes, we will consider almost complex structures satisfying a slightly more generic condition $(J'4')$, which is simply the following:

Definition 9.8. *We say an almost complex structure J on $\Sigma \times \Delta$ satisfies condition $(J'4')$ if there is a 2-place distribution ξ on $\Delta \times \Sigma$ such that the restriction of ω to ξ is nondegenerate, J preserves ξ, and the restriction of J to ξ is compatible with ω. Also ξ is tangent to Σ near $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma}) \times \Delta$ and near $\Sigma \times \partial\Delta$. We also assume that the cylindrical almost complex structures on the ends of $\Sigma \times \Delta$ satisfy $(J5')$.*

We now proceed with a Lemma to count holomorphic triangles in the diagram $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ satisfying a matching condition. Suppose that $\psi \in \pi_2(a, x, b)$ is a homology triangle and $\mathbf{d} \in \mathrm{Sym}^k(\Delta)$ is a divisor. We define

$$\mathcal{M}(\psi, \mathbf{d}) = \{u \in \mathcal{M}(\psi) : \rho^p(u) = \mathbf{d}\}$$

and similarly

$$\mathcal{M}_{(a,x,b)}(\mathbf{d}) = \coprod_{\substack{\psi \in \pi_2(a,x,b) \\ \mu(\psi) = 2k}} \mathcal{M}(\psi, \mathbf{d}).$$

Finally we define

$$M(\mathbf{d}) = \#\mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}).$$

Lemma 9.9. *If $\psi \in \pi_2(a, \theta^+, b)$ then with the multiplicities m_1, m_2, m_3 shown in Figure 14, we have*

$$\mu(\psi) = 2m_1 = m_2 + m_3.$$

If $\psi \in \pi_2(a, \theta^-, b)$ then

$$\mu(\psi) = 2m_1 - 1 = m_2 + m_3.$$

Proof. As usual we verify that the formulas hold for the "trivial" Maslov index zero triangle $\psi_0 \in \pi_2(a, \theta^+, b)$ and then show that the above formula respects splicing in elements of $\pi_2(s, t)$ for various s and t. □

Lemma 9.10. *For generic $\mathbf{d} \in \mathrm{Sym}^k(\Delta)$ and for generic almost complex structures on $\Sigma \times \Delta$ satisfying $(J'4')$, we have that*

$$M(\mathbf{d}) \equiv 1 (\mathrm{mod}\, 2).$$

Proof. This proceeds very similarly to the proof of Lemma 6.19. We first show that $M(\mathbf{d})$ is independent of $\mathbf{d}$. To do this we pick a path $\mathbf{d}_t$ between $\mathbf{d}_0$ and $\mathbf{d}_1$. We now consider the space

$$\coprod_t \mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}_t).$$

First note that the almost complex structure $(J'4')$ prohibits closed surfaces mapped into $(\mathrm{int}\, \Delta)$ from appearing. Hence the only kinds of degenerations which can appear as ends correspond to strip breaking and boundary degenerations, as well as setting $t = 0$ or 1. The first two types of ends correspond to

$$\mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}_i)$$

for $i = 0, 1$. The other two correspond to degenerations of a family of holomorphic triangles into a broken holomorphic triangle. Using the Maslov index in Lemma 9.9 and the fact that the triangular component must satisfy a matching condition with a $\mathbf{d}_t$, we easily see that the only type of degenerations that can occur corresponds to the following space:

$$\coprod_{\substack{t \in I \\ \psi \in \pi_2(a,\theta^-,b) \\ \phi \in \pi_2(\theta^+,\theta^-) \\ \mu(\psi) = 2k-1, \mu(\phi) = 1}} \widehat{\mathcal{M}}(\phi) \times \mathcal{M}(\psi, \mathbf{d}).$$

Since the complex $\widehat{CF}(\mathbb{T}^2, \beta_0, \gamma_0, w)$ has vanishing differential (modulo 2), the above space has an even number of ends, and hence contributes nothing to the total count of ends. We thus conclude that $M(\mathbf{d}_0) - M(\mathbf{d}_1)$ is even, as we wanted.

We now wish to show that $M(\mathbf{d}) \equiv 1$ for appropriate choice of $\mathbf{d}$. To prove this, let $\mathbf{d}_T$ be a divisor with k points spaced at least T apart in Δ, such that as $T \to \infty$, the points in $\mathbf{d}_T$ each approach $v_{\alpha\beta} \subseteq \Delta$. In the cylindrical end of Δ limiting to $v_{\alpha\beta}$, we assume that the I-component of each point in $\mathbf{d}_T$ approaches a value in $(0,1)$. A sequence of curves in $\mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}_T)$ limits to a collection of k Maslov index 2 curves u_i representing classes in $\pi_2(a,a)$ each of which satisfies a matching condition $\rho^p(u_i) = d_i$ for some $d_i \in I \times \mathbb{R}$, as well as a triangular component τ with $[\tau] \in \pi_2(a, \theta^+, b)$ with $\mu([\tau]) = 0$. The triangular component τ is easily seen to be restricted to the "trivial" triangular domain of Maslov index zero, which has a unique representative. Using the proof of [Lip06, Prop. A.3], for almost complex structures on $\Sigma \times \mathbb{D}$ satisfying $(J5')$, we know that the space of holomorphic curves in $\pi_2(a,a)$ with Maslov index 2 satisfying the matching condition with d_i has a unique element. Thus by gluing, we conclude that for sufficiently large T, $M(\mathbf{d}_T) \equiv 1 (\mathrm{mod}\, 2)$, from which the result follows. □

We are now in position to prove Theorem 9.7.

Proof of Theorem 9.7. We first consider almost complex structures satisfying $(J'4')$. Suppose that $\psi = \psi_1 \# \psi_2 \in \pi_2(\mathbf{x} \times a, \mathbf{y} \times \theta^+, \mathbf{z} \times b)$ and suppose that ψ is a Maslov index zero homology triangle which has representatives v_T for arbitrarily large neck length T. Using the combinatorial Maslov index as well as Lemma 9.9, we have that

$$\mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1).$$

Thus we conclude that $\mu(\psi_1) = 0$. Letting $T \to \infty$ we can extract a limit of v_T as a pair of broken triangles V_1 on $\mathcal{T}$ and V_2 on $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$. Since $\mu(\psi_1) = 0$, we conclude that V_1 cannot contain any flowlines, boundary degenerations, or closed surfaces, and hence must be represented by a single triangle v_1. Now in V_2 there must be a matching component, i.e. a triangular component v_2 with

$$\rho^q(v_1) = \rho^p(v_2).$$

We now claim that there are no other components of v_2. The remaining components are represented by flowlines, boundary degenerations, and closed surfaces, they must each have positive Maslov index if they are nonconstant. But since the remaining components have total Maslov index zero, they must be constant. Hence we know that for all sufficiently large T, v_T is obtained by gluing a preglued triangle (v_1, v_2).

Conversely given a holomorphic triangle v_1 on $\mathcal{T}$, according to Lemma 9.10 there is one (modulo 2) holomorphic triangle v_2 on $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ satisfying

$$\rho^q(v) = \rho^p(v_2).$$

Hence we know from gluing that for sufficiently large T, if $\psi_1 \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$, the total count of $\mathcal{M}(\psi_1)$ is equal to the total count of

$$\coprod_{\substack{\psi_2 \in \pi_2(a, \theta^+, b) \\ n_p(\psi_2) = n_q(\psi_1)}} \mathcal{M}(\psi_1 \# \psi_2).$$

□

We need one additional gluing lemma. Theorem 9.7 shows that the triangle maps are invariant under $(1,2)$-stabilization. It also will be used to show that 2-handle maps cancel 1-handle maps. We need the following lemma to show that 3-handle maps cancel 2-handle maps.

Theorem 9.11. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is strongly $\mathfrak{s}$-admissible and let $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ be the Heegaard triple similar to Figure 14 with α_0 and γ_0 isotopic curves with $|\alpha_0 \cap \beta_0| = |\beta_0 \cap \gamma_0| = 1$ and $|\alpha_0 \cap \gamma_0| = 2$. Denote $\overline{\mathcal{T}} = (\Sigma \# \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w})$. Suppose that $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma}})$ and that $\overline{\mathfrak{s}}$ denotes the corresponding* Spin^c *structure in* $\mathrm{Spin}^c(X_{\overline{\boldsymbol{\alpha}}\overline{\boldsymbol{\beta}}\overline{\boldsymbol{\gamma}}})$. *Then for sufficiently stretched and generic almost complex structure, there is an identification of the $\mathbf{z} \times \theta^-$ component of*

$$F^\circ_{\overline{\mathcal{T}}, \overline{\mathfrak{s}}}((\mathbf{x} \times a) \otimes (\mathbf{y} \times b))$$

and the $\mathbf{z}$ component of

$$F^\circ_{\mathcal{T}, \mathfrak{s}}(\mathbf{x} \otimes \mathbf{y}).$$

The proof is identical to the proof of Theorem 9.7, so we omit it.

9.3. **Invariance of maps associated to a framed link.** We are now in position to prove invariance of the maps associated to a framed link. We first show invariance of the map for a fixed bouquet, and then show invariance from the bouquet.

Lemma 9.12. *Suppose we are given a framed link $\mathbb{L} \subseteq Y$ and a bouquet $B(\mathbb{L})$, and two Heegaard triples $\mathcal{T}_i = (\Sigma_i, \boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, \mathbf{w})$ for $i = 1, 2$, which are both strongly $\mathfrak{s}$-admissible for an $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$ and subordinate to $B(\mathbb{L})$. Let $\mathcal{H}_i$ denote the diagram $(\Sigma_i, \boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \mathbf{w})$ for Y and let $\overline{\mathcal{H}}_i$ denote the diagram $(\Sigma_i, \boldsymbol{\alpha}_i, \boldsymbol{\gamma}_i, \mathbf{w})$ for $Y(\mathbb{L})$. Then the following diagram commutes:*

$$\begin{array}{ccc} HF^\circ(\mathcal{H}_1, \mathfrak{s}|_Y) & \xrightarrow{F_{Y,\mathbb{L},\mathcal{T}_1,\mathfrak{s}}} & HF^\circ(\overline{\mathcal{H}}_1, \mathfrak{s}|_{Y(\mathbb{L})}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}_1 \to \mathcal{H}_2}} & & \downarrow{\scriptstyle \Phi_{\overline{\mathcal{H}}_1 \to \overline{\mathcal{H}}_2}} \\ HF^\circ(\mathcal{H}_2, \mathfrak{s}|_Y) & \xrightarrow{F_{Y,\mathbb{L},\mathcal{T}_2,\mathfrak{s}}} & HF^\circ(\overline{\mathcal{H}}_2, \mathfrak{s}|_{Y(\mathbb{L})}) \end{array}$$

Before we proceed with the proof of Lemma 9.12, we state the following lemma:

Lemma 9.13. *Suppose that Y is a 3-manifold with a framed link $\mathbb{L}$ and a bouquet $B(\mathbb{L})$. Then there is a Heegaard triple subordinate to $B(\mathbb{L})$ and any two such diagrams can be connected by a sequence of the following moves:*

1. *isotopies and handleslides amongst the $\boldsymbol{\alpha}$ curves;*
2. *isotopies and handleslides amongst the $\beta_{n+1}, \dots, \beta_{g+\ell-1}$, while performing the the same move to the corresponding $\gamma_{n+1}, \dots, \gamma_{g+\ell-1}$ as well;*
3. *stabilizations and destabilizations;*
4. *for $i \in \{1, \dots, n\}$ an isotopy of β_i, or a handleslide of β_i across a β_j with $j \in \{n+1, \dots, g+\ell-1\}$;*
5. *for $i \in \{1, \dots, n\}$ an isotopy of γ_i or a handleslide of γ_i across a γ_j with $j \in \{n+1, \dots, g+\ell-1\}$*
6. *a diffeomorphism isotopic to the identity in Y relative $N(B(\mathbb{L})) \cup \mathbf{w}$.*

Proof. This follows immediately from [Juh15, Lem. 6.5], where the analogous statement is proven in the sutured setting. □

We are now in position to prove that the maps $F^\circ_{\mathbb{L},\mathcal{T},\mathfrak{s}}$, are invariant of the Heegaard diagram subordinate to B:

Proof of Lemma 9.12. It is sufficient to show that the maps $F^\circ_{\mathbb{L},\mathcal{T},\mathfrak{s}}$ commute with the moves listed in Lemma 9.13.

First note that invariance under move (6) is immediate since an isotopy which fixes $N(B(\mathbb{L})) \cup \mathbf{w}$ induces an isotopy of the Heegaard surface Σ in both Y and $Y(\mathbb{L})$ making the obvious diagram commute.

Invariance under move (3) follows immediately from Theorem 9.7.

Invariance under moves (1),(2),(4) and (5) follow identically to [Juh15, Thm. 6.9], the proof in the sutured context. This essentially follows from associativity of the triangle maps, as well as invariance of the top degree generator of $HF^-((S^1 \times S^2)^{\#n}, \mathfrak{s}_0)$. □

Given a framed link $\mathbb{L}$ and a bouquet $B(\mathbb{L})$, we let $F^\circ_{Y,\mathbb{L},\mathfrak{s},B}$ denote $F^\circ_{Y,\mathbb{L},\mathfrak{s},\mathcal{T}}$ for any triple $\mathcal{T}$ subordinate to B.

Lemma 9.14. *If B and B' are two different bouquets for a framed link $\mathbb{L} \subseteq Y$, then the maps $F^\circ_{\mathbb{L},\mathfrak{s},B}$ and $F^\circ_{\mathbb{L},\mathfrak{s},B'}$ agree.*

Proof. The proof is identical to the proof of the statement in the sutured context (the end of the proof of [Juh15, Thm. 6.9]). We can move from one bouquet to another by a sequence of small isotopies of edges of the bouquet, and replacing one edge at a time. If we want to replace an edge e with another edge e' connected to the same link component, if e and e' are disjoint and B denotes the bouquet with edge e and B' denotes the bouquet with edge e', then we can find two Heegaard triples $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ and $\mathcal{T}' = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$ subordinate to B and B' respectively, such that $\boldsymbol{\beta}'$ is obtained from $\boldsymbol{\beta}$ by a sequence of handleslides and isotopies, and $\boldsymbol{\gamma}'$ is obtained from $\boldsymbol{\gamma}$ by a sequence of handleslides and isotopies. The

commutativity of the appropriate diagram then follows as in proof of invariance of the maps $F^\circ_{\mathbb{L},\mathfrak{s},\mathcal{T}}$ under moves (1),(2), (4) and (5). □

We write simply $F^\circ_{Y,\mathbb{L},\mathfrak{s}}$ for the map $F^\circ_{Y,\mathbb{L},\mathfrak{s},B}$ for any bouquet B.

Lemma 9.15. *If $\mathbb{L}_1$ and $\mathbb{L}_2$ are two disjoint, framed links in Y, then*

$$F^\circ_{Y(\mathbb{L}_1),\mathbb{L}_2,\mathfrak{s}_2} \circ F^\circ_{Y,\mathbb{L}_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W(\mathbb{L}_1)(\mathbb{L}_2)) \\ \mathfrak{s}|_{W(\mathbb{L}_1)}=\mathfrak{s}_1, \mathfrak{s}|_{W(\mathbb{L}_2)}=\mathfrak{s}_2}} F^\circ_{Y,\mathbb{L}_1\cup\mathbb{L}_2,\mathfrak{s}}.$$

Proof. This follows as in [OS06, Prop. 4.9] from associativity of the triangle maps [OS04b, Thm. 8.16]. □

10. Invariance of path cobordism maps

In this section we prove invariance of cobordism maps associated to path cobordisms. This will follow from the framework established in [Juh15] and [Juh14], as well as the gluing lemmas we've proven already. For our purposes, the main result we need is [Juh15, Thm. 8.9]. Juhász phrases these results in terms of cobordisms of sutured manifolds. Given a multibased 3-manifold $(Y, \mathbf{w})$, we can remove regular neighborhoods of the basepoints and add a single suture to each resulting boundary component to get a sutured manifold $Y(\mathbf{w})$. By removing a regular neighborhood of the paths in our cobordism, we get a "special cobordism" and hence can apply Juhász' result directly. A special cobordism is a sutured cobordism from (M, γ) to (M', γ'), where the Z part of the boundary is of the form $\partial M \times I$. On such a cobordism, one can put the unique tight contact structure induced by the sutures on ∂M, so the contact geometry plays no interesting role. It is important to note that if $(W, \boldsymbol{\gamma})$ is a path cobordism, and f is a Morse function on W which has $\boldsymbol{\gamma}$ as flowlines, then f naturally induces a Morse function on the sutured cobordism obtained by removing a regular neighborhood of $\boldsymbol{\gamma}$ from W.

According to [Juh15, Thm. 8.9], in order to define maps associated to special cobordisms, it is sufficient to define handle attachment maps

$$F^\circ_{Y,\mathbb{S},\mathfrak{s}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}|_Y) \to HF^\circ(Y(\mathbb{S}), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S})})$$

for $\mathbb{S}$ a framed 0-sphere, link, or 2-sphere and $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}))$, such that 7 axioms are satisfied. We list the axioms for the convenience of the reader. We have rephrased the axioms slightly since Juhász considers sutured manifolds in full generality, but we need his result only for sutured manifolds of the form $Y(\mathbf{w})$ for a closed $(Y^3, \mathbf{w})$.

Theorem 10.1 ([Juh15, Thm. 8.9])**.** *To define a functorial cobordism maps*

$$F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} : HF^\circ(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathbf{w}_2, \mathfrak{s}|_{Y_2})$$

satisfying the standard Spin^c *composition law for path cobordisms $(W, \boldsymbol{\gamma}, \mathfrak{s}) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$, it is sufficient to define for each framed 0-sphere, link, or 2-sphere $\mathbb{S} \subseteq Y \setminus \mathbf{w}$ and* Spin^c *structure $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}))$ a map*

$$F^\circ_{Y,\mathbb{S},\mathfrak{s}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}|_Y) \to HF^\circ(Y(\mathbb{S}), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S})})$$

which satisfy the following axioms:

1. *If $d \in \mathrm{Diff}_0(Y, \mathbf{w})$, then $d_* : HF^\circ(Y, \mathbf{w}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w}, \mathfrak{s})$ is the identity.*
2. *Consider a diffeomorphism $d : (Y, \mathbf{w}) \to (Y', \mathbf{w}')$ between closed 3-manifolds, and let $\mathbb{S} \subseteq Y \setminus \mathbf{w}$ be a framed 0-sphere, link, or 2-sphere and let $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}))$. Let $\mathbb{S}' = d(\mathbb{S})$. Let*

$$d^{\mathbb{S}} : (Y(\mathbb{S}), \mathbf{w}) \to (Y'(\mathbb{S}'), \mathbf{w}') \text{ and } D^{\mathbb{S}} : W(\mathbb{S}) \to W(\mathbb{S}')$$

be the induced diffeomorphisms, and write $\mathfrak{s}' = D^{\mathbb{S}}(\mathfrak{s}) \in \mathrm{Spin}^c(W(\mathbb{S}'))$. Then the following diagram is commutative:

$$\begin{array}{ccc} HF^\circ(Y, \mathbf{w}, \mathfrak{s}|_Y) & \xrightarrow{F^\circ_{Y,\mathbb{S},\mathfrak{s}}} & HF^\circ(Y(\mathbb{S}), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S})}) \\ \downarrow{\scriptstyle d_*} & & \downarrow{\scriptstyle (d^{\mathbb{S}})_*} \\ HF^\circ(Y', \mathbf{w}', \mathfrak{s}'|_{Y'}) & \xrightarrow{F^\circ_{Y',\mathbb{S}',\mathfrak{s}'}} & HF^\circ(Y'(\mathbb{S}'), \mathbf{w}', \mathfrak{s}'|_{Y'(\mathbb{S}')}) \end{array}.$$

(3) *If $(Y,\mathbf{w})$ is a 3-manifold and $\mathbb{S}$ and $\mathbb{S}'$ are disjoint framed 0-sphere, link, or 2-sphere in Y then $Y(\mathbb{S})(\mathbb{S}') = Y(\mathbb{S}')(\mathbb{S})$ and we denote this manifold by $Y(\mathbb{S},\mathbb{S}')$. There is an isotopically unique equivalence*

$$\mathcal{W} \stackrel{\text{def}}{=} W(Y(\mathbb{S}),\mathbb{S}') \circ W(Y,\mathbb{S}) \cong W(Y(\mathbb{S}'),\mathbb{S}) \circ W(Y,\mathbb{S}').$$

Given $\mathfrak{s} \in \mathrm{Spin}^c(\mathcal{W})$, this equivalence allows us to define $\mathfrak{s}_{\mathbb{S}} = \mathfrak{s}|_{W(Y,\mathbb{S})}$, $\mathfrak{s}_{\mathbb{S}'} = \mathfrak{s}|_{W(\mathbb{S}')}$, $\mathfrak{s}_{\mathbb{S},\mathbb{S}'} = \mathfrak{s}|_{W(Y(\mathbb{S}),\mathbb{S}')}$ and $\mathfrak{s}_{\mathbb{S}',\mathbb{S}} = \mathfrak{s}|_{W(Y(\mathbb{S}'),\mathbb{S})}$. If at least one of $\mathbb{S}$ and $\mathbb{S}'$ is not 1-dimensional, then the following diagram is commutative:

$$\begin{array}{ccc} HF^\circ(Y,\mathbf{w},\mathfrak{s}|_Y) & \xrightarrow{F^\circ_{Y,\mathbb{S},\mathfrak{s}|_{\mathbb{S}}}} & HF^\circ(Y(\mathbb{S}),\mathbf{w},\mathfrak{s}|_{Y(\mathbb{S})}) \\ \downarrow{\scriptstyle F^\circ_{Y,\mathbb{S}',\mathfrak{s}|_{\mathbb{S}'}}} & & \downarrow{\scriptstyle F^\circ_{Y(\mathbb{S}),\mathbb{S}',\mathfrak{s}|_{\mathbb{S},\mathbb{S}'}}} \\ HF^\circ(Y(\mathbb{S}'),\mathbf{w},\mathfrak{s}|_{Y(\mathbb{S}')}) & \xrightarrow{F^\circ_{Y(\mathbb{S}'),\mathbb{S},\mathfrak{s}|_{\mathbb{S}',\mathbb{S}}}} & HF^\circ(Y(\mathbb{S},\mathbb{S}'),\mathbf{w},\mathfrak{s}|_{Y(\mathbb{S},\mathbb{S}')}). \end{array}$$

(4) *Let $\mathbb{L} = \mathbb{L}_1 \cup \mathbb{L}_2$ be a parition of a framed link in Y and write $\mathcal{W}_1 = W(Y,\mathbb{L}_1), \mathcal{W}_2 = W(Y(\mathbb{L}_1),\mathbb{L}_1)$ and $\mathcal{W} = \mathcal{W}_2 \circ \mathcal{W}_1 \cong W(\mathbb{L})$. If $\mathfrak{s}_i \in \mathrm{Spin}^c(\mathcal{W}_i)$ for $i = 1,2$ then*

$$F^\circ_{Y(\mathbb{L}_1),\mathbb{L}_2,\mathfrak{s}_2} \circ F^\circ_{Y,\mathbb{L}_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in\mathrm{Spin}^c(\mathcal{W}) \\ \mathfrak{s}|_{\mathcal{W}_2}=\mathfrak{s}_2 \\ \mathfrak{s}|_{\mathcal{W}_1}=\mathfrak{s}_1}} F^\circ_{Y,\mathbb{L},\mathfrak{s}}.$$

(5) *If $\mathbb{S}' \subseteq Y(\mathbb{S})$ intersects the belt sphere of the handle attached along $\mathbb{S}$ once transversely, then there is a diffeomorphism $\phi : Y \to Y(\mathbb{S})(\mathbb{S}')$ (which is the identity on $Y \cap Y(\mathbb{S})(\mathbb{S}')$ and is unique up to isotopy), such that for every $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}') \circ W(\mathbb{S}))$, if we write $\mathfrak{s}|_{\mathbb{S}} = \mathfrak{s}|_{W(\mathbb{S})}$ and $\mathfrak{s}|_{\mathbb{S}'} = \mathfrak{s}|_{W(\mathbb{S}')}$, then $\mathfrak{s}|_{Y(\mathbb{S})(\mathbb{S}')} = \phi_*(\mathfrak{s}|_Y)$ and*

$$F^\circ_{Y(\mathbb{S}),\mathbb{S}',\mathfrak{s}_{\mathbb{S}'}} \circ F^\circ_{Y,\mathbb{S},\mathfrak{s}_{\mathbb{S}}} = \phi_*.$$

(6) *Suppose that $\mathbb{L}$ is a framed link in $Y \setminus \mathbf{w}$ and $\mathbb{L}'$ is obtained by sliding a component L_i over L_j along a framed arc a. Let $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$ and let $\mathfrak{s}' = (D_1^{L_1})_*(\mathfrak{s})$, $\mathfrak{t} = \mathfrak{s}|_Y, \mathfrak{t}_{\mathbb{L}} = \mathfrak{s}|_{Y(\mathbb{L})}, \mathfrak{t}_{\mathbb{L}'} = \mathfrak{s}'|_{Y(\mathbb{L}')}$, where $D_1^{L_1} : W(\mathbb{L}) \to W(\mathbb{L}')$ is the diffeomorphism described just prior to Theorem 8.9 in* [Juh15]. *If $d_1^{L_i} : Y(\mathbb{L}) \to Y(\mathbb{L}')$ is the restriction of $D_1^{L_i}$, then the following diagram is commutative:*

$$\begin{array}{ccc} HF^\circ(Y,\mathbf{w},\mathfrak{t}) & \xrightarrow{\mathrm{Id}} & HF^\circ(Y,\mathbf{w},\mathfrak{t}) \\ \downarrow{\scriptstyle F^\circ_{Y,\mathbb{L},\mathfrak{s}}} & & \downarrow{\scriptstyle F^\circ_{Y,\mathbb{L}',\mathfrak{s}'}} \\ HF^\circ(Y(\mathbb{L}),\mathbf{w},\mathfrak{t}_{\mathbb{L}}) & \xrightarrow{(d_1^{L_i})_*} & HF^\circ(Y(\mathbb{L}'),\mathbf{w},\mathfrak{t}_{\mathbb{L}'}) \end{array}.$$

(7) *For every framed 0-sphere, link, or 2-sphere $\mathbb{S}$ and $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S})) = \mathrm{Spin}^c(W(\mathbb{S}))$, we have $F^\circ_{Y,\mathbb{S},\mathfrak{s}} = F^\circ_{Y,\overline{\mathbb{S}},\mathfrak{s}}$ where $\overline{\mathbb{S}}$ denotes the sphere $\mathbb{S}$ with the opposite orientation.*

The axioms are verified in [Juh15] for sutured Floer homology, and hence as a special case, for $\widehat{HF}$. We will verify them in this section for HF^-, HF^+ and HF^∞. Much of the work in verifying the axioms for these flavors follows analogously to the case of $\widehat{HF}$. For sutured Floer homology, since disks and triangles aren't allowed to go near the boundary of the sutured Heegaard surface, one can require 1-handles and stabilizations to occur only in a neighborhood of the boundary, eliminating the need for some of the gluing lemmas we have proven. Our main contributions are the various gluing arguments needed to show relations between maps defined by counting holomorphic curves when curves are allowed to go over the basepoints.

Theorem 10.2. *The maps $F^\circ_{Y,\mathbb{S},\mathfrak{s}}$ defined previously for framed 1- and 3- handles, as well as the maps $F^\circ_{Y,\mathbb{L},\mathfrak{s}}$ for framed links satisfy axioms (1)-(7) of* [Juh15, Thm. 8.9]. *In particular, Theorem A follows.*

Proof. Axioms (1) and (2) follow from the invariance of the Heegaard Floer homology groups, as well as the invariance of the 1-,2- and 3-handle maps. Axiom (3) is immediate if neither framed sphere is 1-dimensional. In the case that one of the $\mathbb{S}$ is a framed link, the diagram in axiom (3) commutes by Theorem 8.8. Axiom (4) follows from Lemma 9.15. Axiom (6), invariance under handleslides in a link $\mathbb{L}$, follows without change as the proof in the sutured setting in [Juh15, Thm. 8.11]. Reversing the orientation of a framed link or sphere has no effect on our maps, so axiom (7) is satisfied.

The remaining axiom, axiom (5), requires more work in our setting. We will consider only the case of canceling 1- and 2-handles in detail, since the case of 2- and 3-handles is formally the same. The work in [Juh15, Thm. 8.11] is not quite sufficient since in that context disks and triangles cannot go over the basepoints, so weaker gluing arguments can be used in that context.

Suppose that $\mathbb{S}_0 \subseteq Y \setminus \mathbf{w}$ is a framed 0-sphere and let $\mathbb{S}_2 \subseteq Y(\mathbb{S}_0)$ denote the canonical 2-sphere in $Y(\mathbb{S}_0)$. Suppose $\mathbb{K} \subseteq Y(\mathbb{S}_0) \setminus \mathbf{w}$ is a framed knot transverse to $\mathbb{S}_2$ with $|\mathbb{K} \cap \mathbb{S}_2| = 1$. According to [Juh15, Thm. 8.11.] there is a diffeomorphism $\phi : Y \to Y(\mathbb{S}_0)(\mathbb{K})$ which is unique up to isotopy fixing $\mathbf{w}$. Suppose that $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}_0)(\mathbb{K}))$. Axiom (5) states that the following relation should hold:

$$F^\circ_{Y(\mathbb{S}_0),\mathbb{K},\mathfrak{s}|_{W(\mathbb{K})}} \circ F^\circ_{Y,\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}} = \phi_*.$$

Following [Juh15, Lem. 8.3], we can pick a bouquet $B(\mathbb{K})$, a diagram

$$\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$$

and a strongly admissible triple

$$\mathcal{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w}) = ((\Sigma \setminus B_1 \cup B_2) \cup A, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w})$$

such that

(1) $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a strongly admissible Heegaard diagram of $(Y, \mathbf{w})$ subordinate to $\mathbb{S}_0$ (i.e. $\mathbb{S}_0 \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w})$);
(2) B_1 and B_2 are two small balls in Σ centered around the points of $\mathbb{S}_0 \subseteq \Sigma$ which don't intersect $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w}$;
(3) A is an annulus in $Y(\mathbb{S}_0)$ such that $\partial A = \partial B_1 \cup \partial B_2$;
(4) the curve α_0 is a homotopically nontrivial curve in A and β_0 is a small Hamiltonian isotopy of α_0 with $|\alpha_0 \cap \beta_0| = 2$;
(5) the curves $\boldsymbol{\gamma}$ are small Hamiltonian isotopies of the $\boldsymbol{\beta}$ curves which intersect $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ transversely;
(6) γ_0 intersects α_0 and β_0 each transversally at a single point and γ_0 intersects no other $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves;
(7) $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$ is a triple subordinate to the bouquet $B(\mathbb{K})$ and β_0 is a meridian of K, and γ_0 is the framing of K;

Now $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\gamma}', \mathbf{w})$ is a diagram for $(Y(\mathbb{S}_0)(\mathbb{K}), \mathbf{w})$ which is the $(1,2)$-stabilization of $\phi_*\mathcal{H}$. Diagrammatically the situation is depicted in Figure 15.

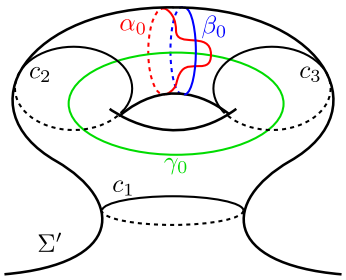

FIGURE 15. The Heegaard triple $\mathbb{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$. The circles c_1, c_2 and c_3 are circles on which we need to stretch to compute the 1-handle map or to apply the stabilization lemma for triangle maps.

We need to show that

$$F^\circ_{\mathbb{K},\mathfrak{s}|_{W(\mathbb{K})}} \circ F^\circ_{\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}} = \phi_*.$$

Let $a = \alpha_0 \cap \gamma_0$ and $b = \beta_0 \cap \gamma_0$. In terms of diagrams, the desired equality becomes

$$F^\circ_{\mathcal{T}',\mathfrak{s}|_{W(\mathbb{K})}}(F^\circ_{\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}}(\mathbf{x}) \otimes (\Theta^+ \times b)) = \Phi_{(\phi_*\mathcal{H}) \to \mathcal{H}'} \circ \phi_*.$$

By stretching along c_1 sufficiently, the map $\Phi_{\mathcal{H} \to \mathcal{H}'}$ is the stabilization isomorphism $\mathbf{z} \mapsto \mathbf{z} \times a$. For almost complex structures which are sufficiently stretched along c_2 and c_3, the 1-handle map $F_{\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}}$ takes the

form $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$ where θ^+ is the higher degree point in $\alpha_0 \cap \beta_0$. Also for almost complex structure sufficiently stretched along c_1, by Lemma 9.7, we have an identification of

$$F^{\circ}_{\mathcal{T}', \mathfrak{s}|_{W(\mathbb{K})}}((\mathbf{x} \times \theta^+) \otimes (\Theta^+ \times b)) = \sigma \circ \phi_*,$$

where σ is the stabilization isomorphism. Thus in particular, the desired equality will be satisfied if there is some almost complex structure such that the 1-handle map takes the form $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$, and such the stabilization map takes the form $\mathbf{y} \mapsto \mathbf{y} \times b$ and the 2-handle maps take form in Theorem 9.7.

To this end, we will need the following subclaim: if J_s is an almost complex structure on Σ' which is split in neighborhoods of c_1, c_2 and c_3, then if we let $J_s(T_1, T_2, T_3)$ denote the almost complex structure obtained by inserting neck lengths of T_1, T_2 and T_3, then there is an N such that if $T_i, T_i' > N$ for $i \in \{1, 2, 3\}$, then the maps

$$\Phi_{J_s(T_1,T_2,T_3) \to J_s(T_1',T_2',T_3')} : CF^{\circ}(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S}_0)}) \to CF^{\circ}(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S}_0)})$$

are $\mathbb{Z}_2[U_{\mathbf{w}}]$-chain homotopic to the map

$$\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}$$

.

Before we prove the subclaim, let's see how it proves the desired result. If we pick such an N, and if $T_2, T_3 > N$, then the almost complex structure $J_s(T, T_2, T_3)$ is "long enough to compute the 1-handle map" for any $T > N$. Hence for any T, the 1-handle map takes the form $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$ with respect to the almost complex structure $J_s(T, T_2, T_3)$. By Theorem 9.7 to the triple $\mathcal{T}' = (\Sigma', \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w})$, we know that for sufficiently large T, we have

$$F^{\circ}_{\mathcal{T}', \mathfrak{s}|_{W(\mathbb{K})}}((\mathbf{x} \times \theta^+) \otimes (\Theta^+ \times b)) = \sigma \circ \phi_*,$$

implying by our previous observations that the following equality is satisfied for a particular choice of diagrams

$$F^{\circ}_{\mathbb{K}, \mathfrak{s}|_{W(\mathbb{K})}} \circ F^{\circ}_{\mathbb{S}_0, \mathfrak{s}|_{W(\mathbb{S}_0)}} = \phi_*.$$

We now prove the subclaim. To see this pick a sequence of 6-tuples of the form $(T_{1,n}, T_{2,n}, T_{3,n}, T_{1,n}', T_{2,n}', T_{3,n}')$ where $T_{i,n}$ and $T_{i,n}'$ all approach $+\infty$. Pick an interpolating almost complex structure $\widehat{J}_n$ which is constant outside of a neighborhood of the circles c_i. Suppose that u_n is a sequence of $\widehat{J}_n$-holomorphic strips which represent a Maslov index zero homology class. The almost complex structure $\widehat{J}_n$ converge to an almost complex structure $\widehat{J}$ on $\Sigma_\infty = \Sigma \cup_{c_1} S^2 \cup_{c_2, c_3} S^2$ which is topologically obtained by collapsing the circles c_i to points. Suppose that all u_n represent the homology class $\phi = \phi_1 \# \phi_2 \# \phi_3$. Taking a Gromov limit of u_n, we get holomorphic representatives of each ϕ_i by broken holomorphic strips U_i. If $\phi \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^+)$ or $\pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^-)$, then by Lemma 6.3 we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) + \mu(\phi_3) - 6m = \mu(\phi_1),$$

where m is the multiplicity in the region containing the circles c_1, c_2 and c_3. Hence if $\mu(\phi) = 0$, then $\mu(\phi_1) = 0$. Since U_1 is a broken holomorphic strip on an $\mathbb{R}$-invariant almost complex structure on Σ with Maslov index zero, we conclude that U_1 can only have constant components. Hence the diagonal entries of the transition matrix are equal to the identity. Similarly if $\phi \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^-)$ has Maslov index zero, then by Lemma 6.3, we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) + \mu(\phi_3) - 6m = \mu(\phi_1) + 1.$$

If ϕ has $\widehat{J}_n$-holomorphic representatives for arbitrarily large n, then taking a weak limit yields a broken representative U_1 of ϕ_1, implying that $\mu(\phi_1) \geq 0$. Hence $\mu(\phi) = \mu(\phi_1) + 1 \geq 1$, implying that no such u_n are counted in the transition matrix. Hence the lower left entry of the transition map is zero, completing the proof of the subclaim.

The analogous statement holds for 3-handles which cancel a 2-handle. The only difference is that we must use the gluing result in Theorem 9.11 for the model computation. □

The following is a corollary of the above theorem and the definition of a functor which is split over Spin^c structures ([Juh15, Def. 8.7])

Corollary 10.3. *If $(W,\boldsymbol{\gamma})$ is a path cobordism which decomposes as*

$$(W,\boldsymbol{\gamma}) = (W_2,\boldsymbol{\gamma}_2)\circ(W_1,\boldsymbol{\gamma}_1)$$

then the following composition law holds:

$$F^\circ_{W_2,\boldsymbol{\gamma}_2,\mathfrak{s}_2}\circ F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in\mathrm{Spin}^c(W)\\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}} F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}}.$$

11. Construction and invariance of the graph cobordism maps when $\pi_0(Y_i)\twoheadrightarrow\pi_0(W)$

In this section, we construct the graph cobordism maps using the path cobordism maps, and the graph action map, though in this section we focus only on graph cobordisms $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ where the inclusions $\pi_0(Y_i)\to\pi_0(W)$ are surjective for $i=1,2$. The construction in this section actually works if we assume only that $\pi_0(Y_2)\to\pi_0(W)$ is surjective, but it provides no benefit to work in this intermediate level of generality. Thus we assume in this section that $\pi_0(Y_i)\to\pi_0(W)$ is surjective, and consider the fully general case in Section 13, once we have established the composition law for cobordisms with $\pi_0(Y_i)\twoheadrightarrow\pi_0(W)$.

In this section, we prove the following weaker version of Theorem C:

Theorem 11.1. *Given a graph cobordism $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ where Γ is a ribbon graph with coloring $(\sigma,\mathfrak{P})$, and $\pi_0(Y_i)\twoheadrightarrow\pi_0(W)$ is surjective, there is a homomorphism of $\mathbb{Z}_2[U_{\mathfrak{P}}]$-modules*

$$F_{W,\Gamma,\mathfrak{s}}: HF^\circ(Y_1,\mathbf{w}_1,\sigma|_{\mathbf{w}_1},\mathfrak{s}|_{Y_1})\to HF^\circ(Y_2,\mathbf{w}_2,\sigma|_{\mathbf{w}_2},\mathfrak{s}|_{Y_2}).$$

These maps are invariants of the cobordism with embedded colored ribbon graph cobordism.

11.1. **0- and 4-handle maps.** In this subsection we define the 0- and 4-handle maps. Unlike the 1-,2-, or 3-handle maps, the 0- and 4-handle maps are easy to define.

Let $(\mathbb{T}^2,\alpha_0,\beta_0,w_0)$ be the standard diagram for (S^3,w_0) where $\alpha_0\cap\beta_0=\{c\}$. Given a three manifold $(Y,\mathbf{w})$ and a Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w})$ of $(Y,\mathbf{w})$, the zero handle map

$$F^\circ_{0,\mathfrak{s}}: HF^\circ(Y,\mathbf{w},\mathfrak{s})\to HF^\circ(Y\sqcup S^3,\mathbf{w}\cup\{w_0\},\mathfrak{s}\sqcup\mathfrak{s}_0)$$

is defined on complexes as

$$F^\circ_{0,\mathfrak{s}}: CF^\circ(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathfrak{s})\to CF^\circ(\Sigma\sqcup\mathbb{T}^2,\boldsymbol{\alpha}\cup\{\alpha_0\},\boldsymbol{\beta}\cup\{\beta_0\},\mathbf{w}\cup\{w_0\},\mathfrak{s}\sqcup\mathfrak{s}_0)$$

by the formula

$$F^\circ_{0,\mathfrak{s}}(\mathbf{x}) = \mathbf{x}\times c.$$

The map is natural in the sense that it commutes with change of diagrams maps on Y.

In the opposite direction, given a based 3-manifold $(Y,\mathbf{w})$, there is a 4-handle map

$$F^\circ_{4,\mathfrak{s}}: HF^\circ(Y\sqcup S^3,\mathbf{w}\cup\{w_0\},\mathfrak{s}\sqcup\mathfrak{s}_0)\to HF^\circ(Y,\mathbf{w},\mathfrak{s}),$$

which is defined on complexes as a map

$$F^\circ_{4,\mathfrak{s}}: CF^\circ(\Sigma\sqcup\mathbb{T}^2,\boldsymbol{\alpha}\cup\{\alpha_0\},\boldsymbol{\beta}\cup\{\beta_0\},\mathbf{w}\cup\{w_0\},\mathfrak{s}\sqcup\mathfrak{s}_0)\to CF^\circ(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathfrak{s})$$

by the formula

$$F^\circ_{4,\mathfrak{s}}(\mathbf{x}\times c) = \mathbf{x}.$$

When there is no ambiguity in doing so, we will often omit the Spin^c structure $\mathfrak{s}$ and write only F°_0 or F°_4.

11.2. **Main construction.** We now proceed with the main construction of the graph cobordism maps. For the construction, we need to first fix some data about the graph in a graph cobordism (W,Γ):

- fix a subdivision of Γ such that there are no edges from a vertex to itself (in subsection 11.5 we will show that the graph cobordism maps are invariant under subdivision);
- fix an absolute ordering on the edges incident to each vertex (in subsection 11.6 we will show that the graph cobordism maps are invariant under cyclic reordering of the vertices);

Suppose that W^4 is a compact, oriented 4-manifold with boundary and Γ is an embedded graph. Let $N(\Gamma)$ denote a regular neighborhood of Γ. The graph Γ yields a handle decomposition, $\mathfrak{H}$, of $N(\Gamma)$ in the following sense: to each vertex in $V(\Gamma)$ which is not in ∂W, we associate a 0-handle, i.e. a subset diffeomorphic to a ball B^4 centered at the vertex v. To each edge of Γ, we associate a 1-handle, i.e. a subset diffeomorphic to $I \times B^3$. To a vertex $v \in V(\Gamma) \cap \operatorname{int} W^4$, the 1-handles we associated to edges incident to v are glued to the 0-handle associated to v.

We fix some basic notation about embeddings and isotopies of submanifolds and graphs:

Definition 11.2. *If W and Y are compact manifolds with boundary and $Y \hookrightarrow W$ is an embedding, we say that $Y \subseteq W$ is a **neatly embedded submanifold** if $\partial Y \subseteq \partial W$ and Y is transverse to ∂W. If $\Phi_t : W \times I \to W$ is a 1-parameter family of diffeomorphisms, we say that Φ_t is **neat isotopy** if Φ_t maps ∂W into ∂W. If $\phi_t : Y \to W$ is a 1-parameter family of embeddings of Y into W, we say ϕ_t is a **neat isotopy** if $\phi_t(Y)$ is a neat submanifold for all t. If $\Gamma \subseteq W$ is a graph with a distinguished set of vertices denoted $\partial\Gamma$, we say that Γ is a **neatly embedded graph** if $\partial\Gamma$ consists only of 1-valent vertices in Γ, $\Gamma \cap \partial W = \partial\Gamma$ and Γ intersects ∂W transversely. If Γ_t is a 1-parameter family of graphs in W (i.e. a continuous family of maps $\phi_t : \Gamma \to W$ for a fixed graph Γ, with a fixed set of distinguished boundary vertices, $\partial\Gamma$), we say that Γ_t is a **neat graph isotopy** if the $\phi_t(\Gamma)$ is neatly embedded for all t.*

The graph Γ specifies a graph $\widetilde{\Gamma}$ in $\partial N(\Gamma)$, isotopic to Γ inside of $N(\Gamma)$, which we now describe. The handle decomposition of $N(\Gamma)$ induced by Γ induces a decomposition of $\partial N(\Gamma)$ into pieces as well. More precisely, for each vertex $v \in V(\Gamma) \cap \operatorname{int} W$, there is a subset $M(v)$ of $\partial N(\Gamma)$ defined by

$$M(v) = H(v) \cap \partial N(\Gamma)$$

where $H(v)$ denotes the 0-handle corresponding to the vertex v. Note that if v has valence n, then $\partial H(v) \cap \partial N(\Gamma)$ is diffeomorphic to $S^3 \setminus (n \text{ balls})$.

For an edge e, there is a subset $M(e)$ of $\partial N(\Gamma)$ defined by

$$M(e) = H(e) \cap \partial N(\Gamma),$$

where $H(e)$ is the 1-handle associated to the edge e. Note that $H(e)$ is diffeomorphic to $S^2 \times I$.

The graph $\widetilde{\Gamma}$ in $N(\Gamma)$ can now be described as follows. First arbitrarily pick a point on each of the embedded 2-spheres $S \subseteq \partial N(\Gamma)$ corresponding to the boundaries of the subsets $M(e)$ and $M(v)$. In each subset of the form $M(v)$ for $v \in V(\Gamma) \cap \operatorname{int} W$, we place a vertex $\widetilde{v}$ arbitrarily in the interior of $M(v)$, and connect it with n pairwise disjoint paths, each from $\widetilde{v}$ to the marked points on the boundary of $M(v)$. in each subset of the form $M(e)$ for e an edge, there are two boundary spheres, and we pick a path arbitrarily between the distinguished points on the boundary spheres. Patching these all together, we get a graph $\widetilde{\Gamma} \subseteq \partial N(\Gamma)$. There is, of course, ambiguity in choosing the graph $\widetilde{\Gamma}$, but we have the following:

Lemma 11.3. *Any graph $\widetilde{\Gamma}$ constructed using the above procedure is isotopic to Γ by a neat isotopy $\phi_t : W \to W$ which maps $\widetilde{\Gamma}$ into $N(\Gamma)$ for all t. Also any two $\widetilde{\Gamma}$ and $\widetilde{\Gamma}'$ constructed in this manner are neatly isotopic in $\partial N(\Gamma)$.*

Proof. To see that $\widetilde{\Gamma}$ and Γ are isotopic, we simply first isotope $\widetilde{\Gamma}$ so the vertices are in the same locations as the vertices of Γ. After performing this isotopy to $\widetilde{\Gamma}$, the edges of $\widetilde{\Gamma}$ and Γ are clearly homotopic in $N(\Gamma)$, and hence isotopic since $N(\Gamma)$ is a 4-manifold.

We now show that $\widetilde{\Gamma}$ and $\widetilde{\Gamma}'$ are neatly isotopic in $\partial N(\Gamma)$. This is a generalization of the "lightbulb theorem", i.e. the fact that any knot in $S^2 \times S^1$ which intersects $S^2 \times \{*\}$ transversely at a single point is isotopic to the knot $\{p\} \times S^1$. If $v \in V(\Gamma) \cap \operatorname{int} W$ then the portion of $\widetilde{\Gamma}$ inside of $M(v)$ consists of a vertex $\widetilde{v}$ connected to each boundary component of $M(v)$ by a path. We need to show that any two such graphs are isotopic relative $\partial M(v)$. Clearly any two graphs are homotopic. The homotopy can be taken to be an isotopy except at finitely many points of time, where a crossing change occurs. We can change any crossing via the following trick: given a crossing of two arcs a and a', traveling along a' in one direction leads to a boundary component of $M(v)$. We simply isotope a along a' in this direction, then around the boundary component and then back to its original position, but with the crossing changed. □

We now continue to describe the construction of the graph cobordism maps. Before we define them, we need to define more auxiliary data and one additional intermediate map, corresponding to attaching 0-

and 1-handles according to $\mathfrak{H}$. Suppose that $(W,\Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a graph cobordism and Γ is a cyclically ordered graph. Subdivide Γ further so that the edges connecting to Y_2 have as boundary a 1-valent vertex (in Y_2) and a 2-valent vertex (in $\operatorname{int} W$). Let $\boldsymbol{\gamma}_f$ denote the collection of edges connected to Y_2. Let Γ_0 denote the graph $\Gamma \setminus \boldsymbol{\gamma}_f$. Pick a regular neighborhood $N(\Gamma_0)$ of Γ_0. Following Lemma 11.3, let Φ_t be a neat isotopy of W which is fixed outside of a small neighborhood of $\overline{N(\Gamma_0)}$, which maps Γ_0 to a graph $\widetilde{\Gamma}_0$ in $\partial N(\Gamma_0)$ and such that $\Phi_t|_{Y_1}$ is an isotopy of Y_1 which is fixed outside a small neighborhood of $\overline{N(\Gamma) \cap Y_1} = \overline{N(\mathbf{w}_1)}$. Let $\widetilde{\boldsymbol{\gamma}}_f$ denote the image of $\boldsymbol{\gamma}_f$ under the isotopy Φ_t. Let $\phi = \Phi_1|_{Y_1}$. Most of this data is illustrated in Figure 16.

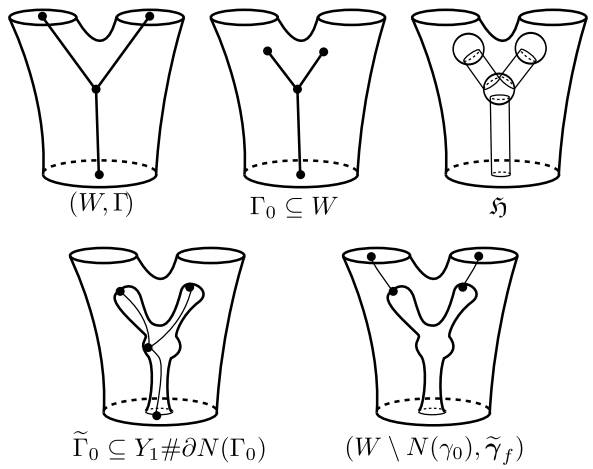

FIGURE 16. The various topological constructions used in the construction of the graph cobordism map.

Suppose that $\mathfrak{t} \in \operatorname{Spin}^c(Y_1)$. We now define a map

$$F^\circ_{\mathfrak{H},\mathfrak{t}} : HF^\circ(Y_1, \phi_*\mathbf{w}, \mathfrak{t}) \to HF^\circ(Y_1 \# \partial N(\Gamma_0), \phi_*\mathbf{w} \cup \mathbf{v}, \mathfrak{t}'),$$

where $\mathbf{v} = \{\widetilde{v} \in V(\widetilde{\Gamma}_0) \cap \operatorname{int} W\}$, which is a composition of 0- and 1-handle maps according to $\mathfrak{H}$. Here $\mathfrak{t}'$ is a Spin^c on $Y_1 \# \partial N(\Gamma_0)$ which is uniquely determined by $\mathfrak{t}$. If $\mathfrak{t} = \mathfrak{s}|_{Y_1}$ for some $\mathfrak{s} \in \operatorname{Spin}^c(W)$, then $\mathfrak{t}' = \mathfrak{s}|_{Y_1 \# \partial N(\Gamma_0)}$. Write $H_1, \dots, H_k$ for the 0-handles of $\mathfrak{H}$ and $H_{k+1}, \dots, H_n$ for the 1-handles. Let $G_i \subseteq \Gamma_0$ denote the union of vertices and edges corresponding to handles $H_1, \dots, H_i$ in the above ordering. If H_i is a zero handle, there is a diffeomorphism $\phi_i : (Y_1 \# \partial N(G_{i-1})) \sqcup S^3 \to Y_1 \# \partial N(G_i)$ which is the identity on the component which is not S^3. Since $MCG(S^3, w_0) = 1$, this diffeomorphism is well defined up to isotopy fixing w_0. If G_i is obtained from G_{i-1} by adding a vertex v_i, we define a map

$$f^\circ_{i,\mathfrak{t}} : HF^\circ(Y_1 \# \partial N(G_{i-1}), \phi_*\mathbf{w} \cup \mathbf{v}^{\le i-1}, \mathfrak{t}) \to HF^\circ(Y_1 \# \partial N(G_i), \phi_*\mathbf{w} \cup \mathbf{v}^{\le i}, \mathfrak{t} \sqcup \mathfrak{t}_0)$$

by

$$f^\circ_{i,\mathfrak{t}} = (\phi_i)_* \circ F^\circ_{0,\mathfrak{t}}$$

where $F^\circ_{0,\mathfrak{t}}$ is the 0-handle map. Here $\mathbf{v}^{\le i} = \mathbf{v}^{\le i-1} \cup \{\widetilde{v}_i\}$ with $\mathbf{v}^{\le 0} = \varnothing$, where $\widetilde{v}_i$ denotes the vertex of $\widetilde{\Gamma}_0$ in H_i. The diffeomorphism ϕ_i is defined up to isotopy, so the map $(\phi_i)_*$ is well defined up to chain homotopy.

If G_i is an edge, we define a map

$$f^\circ_{i,\mathfrak{t}} : HF^\circ(Y_1 \# \partial N(G_{i-1}), \phi_*\mathbf{w} \cup \mathbf{v}, \mathfrak{t}) \to HF^\circ(Y_1 \# \partial N(G_i), \phi_*\mathbf{w} \cup \mathbf{v}, \mathfrak{t}')$$

simply by

$$f^\circ_{i,\mathfrak{t}} = F^\circ_{(Y_1 \# \partial N(G_{i-1}) \times I) \cup H_i, (\phi_*\mathbf{w} \cup \mathbf{v}) \times I, \bar{\mathfrak{t}}},$$

the path cobordism map associated to path cobordism $((Y_1 \# \partial N(G_{i-1}) \times I) \cup H_i, (\phi_*\mathbf{w}) \times I)$. Here $\bar{\mathfrak{t}}$ is the unique Spin^c structure on $(Y_1 \# \partial N(G_{i-1})) \times I \cup H_i$ which restricts to $\mathfrak{t}$ on $Y_1 \# \partial N(G_{i-1})$.

We now define the map $F^\circ_{\mathfrak{H},\mathfrak{t}}$ as the composition

$$F^\circ_{\mathfrak{H},\mathfrak{t}} = \prod_{i=k+1}^{n} f^\circ_{i,\mathfrak{t}_i},$$

where $\mathfrak{t}_1 = \mathfrak{t}$, and $\mathfrak{t}_i$ is either $\mathfrak{t}_{i-1} \sqcup \mathfrak{t}_0$ or $\mathfrak{t}'_{i-1}$, using the above notation, depending on whether G_i is obtained from G_{i-1} by adding a vertex or an edge.

We define the graph cobordism map $F^\circ_{\Gamma,\mathfrak{s}}$ as follows. Write W' for the 4-manifold $W \setminus N(\Gamma_0)$. Write $\mathcal{W}'$ for the path cobordism $\mathcal{W}' = (W', \widetilde{\gamma}_f)$. Notice that W' has boundary equal to $-(Y_1 \# \partial N(\Gamma_0)) \sqcup Y_2$. Define the graph cobordism map as

$$F^\circ_{W,\Gamma,\mathfrak{s}} \stackrel{\text{def}}{=} F^\circ_{\mathcal{W}',\mathfrak{s}|_{W'}} \circ A^V_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} \circ \phi_*. \tag{8}$$

The set V of vertices consists of all vertices in $V(\widetilde{\Gamma}_0)$. Notice that since V consists of all vertices of $V(\widetilde{\Gamma}_0)$, *a-priori* $F^\circ_{W,\Gamma,\mathfrak{s}}$ could depend on the absolute ordering of the edges adjacent to each vertex, since Lemma 7.12 only applies to vertices which are not in the set V. We show in subsection 11.6 that it is invariant under cyclically reordering the edges adjacent to each vertex.

As we proceed to analyze the graph cobordism maps, if we are given a handlebody decomposition $\mathfrak{H}$ of the regular neighborhood of a graph Γ, we will often write

$$F^\circ_{\mathfrak{H}}$$

instead of $F^\circ_{\mathfrak{H},\mathfrak{s}}$ since there is no ambiguity.

Having now defined the graph cobordism maps in terms of some additional choices, we now proceed to show that it is invariant under the additional data and depends only on the Spin^c decorated ribbon-graph cobordism $(W, \Gamma, \mathfrak{s})$.

11.3. **Invariance of the graph cobordism maps from $\widetilde{\Gamma}_0$.** In the construction of the graph cobordism maps, we needed to isotope Γ into $\partial N(\Gamma)$ to get a graph $\widetilde{\Gamma}$. We first need to show that the graph cobordism map is independent of the choice of isotopic graph $\widetilde{\Gamma}_0$ or the choice of isotopy $\Phi_t : W \times I \to I$ sending Γ_0 to $\widetilde{\Gamma}_0$.

Lemma 11.4. *Suppose that $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ are two choices of graphs in $\partial N(\Gamma)$ as constructed in the previous section, and Φ_t and Φ'_t are two isotopies of W mapping Γ_0 to $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ respectively, which are each the identity outside of a neighborhood of $\overline{N(\Gamma_0)}$. Then the graph cobordism maps computed with respect to the pairs $(\widetilde{\Gamma}_0, \Phi_t)$ and $(\widetilde{\Gamma}'_0, \Phi'_t)$ agree.*

Proof. By Lemma 11.3 we can find a neat isotopy $(\widetilde{\Gamma}_0)_t$ between $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ inside of $\partial N(\Gamma_0)$. Such an isotopy of graphs induces an isotopy between $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ inside of $Y_1 \# \partial N(\Gamma_0)$. First extend the isotopy of graphs to an isotopy of $Y_1 \# \partial N(\Gamma_0)$ which is fixed outside of a small neighborhood of $\partial N(\Gamma_0)$. Since this isotopy can be chosen to map $\partial(\partial N(\Gamma_0))$, a collection of 2-spheres, to itself, we can then extend this to a neat isotopy $\Psi_t : W \times I \to W$ which is also fixed outside of a small neighborhood of $\partial N(\Gamma_0)$. Define $\psi : Y_1 \to Y_1$ by $\psi = \Psi_1|_{Y_1}$. Define $\phi = \Phi_1|_{Y_1}$ and $\phi' = \Phi'_1|_{Y_1}$.

To show that the graph cobordism maps computed with respect to the pairs $(\widetilde{\Gamma}_0, \Phi_t)$ and $(\widetilde{\Gamma}'_0, \Phi'_t)$ agree, it is sufficient to show that each square in the following diagram commutes (the diagram has been broken into two lines):

$$\begin{array}{ccccccc}
HF^\circ(Y_1, \mathbf{w}_1) & \xrightarrow{\phi_*} & HF^\circ(Y_1, \widetilde{\mathbf{w}}_1) & \xrightarrow{F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}} & HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}_1) & \xrightarrow{A_{\widetilde{\Gamma}_0}} & HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}_f) \cdots \\
\downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \psi_*} & & \downarrow{\scriptstyle (\Psi_1)_*} & & \downarrow{\scriptstyle (\Psi_1)_*} \\
HF^\circ(Y_1, \mathbf{w}_1) & \xrightarrow{\phi'_*} & HF^\circ(Y_1, \widetilde{\mathbf{w}}'_1) & \xrightarrow{F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}} & HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}'_1) & \xrightarrow{A_{\widetilde{\Gamma}'_0}} & HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}'_f) \cdots
\end{array}$$

$$\begin{array}{ccc} \cdots HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}_f) & \xrightarrow{F^\circ_{(W\setminus N(\Gamma_0),\widetilde{\gamma}_f),\mathfrak{s}|_{W\setminus N(\Gamma_0)}}} & HF^\circ(Y_2, \mathbf{w}_2) \\ \downarrow{\scriptstyle (\Psi_1)_*} & & \downarrow{\scriptstyle (\Psi_1)_* = \mathrm{id}} \\ \cdots HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}'_f) & \xrightarrow{F^\circ_{(W\setminus N(\Gamma_0),\widetilde{\gamma}'_f),\mathfrak{s}|_{W\setminus N(\Gamma_0)}}} & HF^\circ(Y_2, \mathbf{w}_2) \end{array}$$

In the above expression, $\widetilde{\mathbf{w}}_1$ denotes $\widetilde{\Gamma_0} \cap Y_1$ and $\widetilde{\mathbf{w}}_f$ denotes the endpoints of the $\widetilde{\boldsymbol{\gamma}}_f$ arcs which lie in $\operatorname{int} W$.

The composition of the top row is the graph cobordism map computed with respect to $(\widetilde{\Gamma}_0, \Phi_t)$ and the second is the graph cobordism map computed with respect to $(\widetilde{\Gamma}'_0, \Phi'_t)$. All of the squares except for the first one commute simply because of diffeomorphism invariance of the maps involved. Some amount of care is needed in the first square since ϕ' is not necessarily equal to $\psi \circ \phi$. On the other hand, we claim that $\psi \circ \phi$ and ϕ' are isotopic relative $\mathbf{w}_1$. To see this note that ϕ, ψ and ϕ' are all equal to the identity of a neighborhood $N(\mathbf{w}_1) \subseteq Y_1$ of the basepoints. The map $(\phi')^{-1} \circ \psi \circ \phi$ is isotopic to the map resulting from finger moves of the basepoints $\mathbf{w}_1$ around loops which are contained in $N(\mathbf{w}_1)$. Hence $(\phi')^{-1} \circ \psi \circ \phi \simeq \mathrm{id}$ relative $\mathbf{w}$, and hence $\phi' \simeq \psi \circ \phi$ relative $\mathbf{w}$. Hence the diagram commutes, and the graph cobordism map doesn't depend on the choice of $\widetilde{\Gamma}_0$ and Φ_t. □

11.4. **Invariance of the cobordism maps from $N(\Gamma_0)$.** The next step towards proving invariance is show that the graph cobordism maps don't depend on the choice of regular neighborhood $N(\Gamma_0)$. To this end, we note that in the construction of the graph cobordism maps, we created a sort of handle body decomposition of $N(\Gamma_0)$, the combinatorics of which is described by Γ_0. We prove the following lemma about invariance up to isotopy of these handlebody decompositions:

Lemma 11.5. *Suppose that $\Gamma \subseteq W$ is a neatly embedded graph and that $N(\Gamma)$ and $N'(\Gamma)$ are two regular neighborhoods and $\mathfrak{H}$ and $\mathfrak{H}'$ are two handlebody decompositions of the type described in the previous section for $N(\Gamma)$ and $N'(\Gamma)$. Then there is a neat isotopy of W which can be taken to be fixed outside of a given neighborhood of $\overline{N(\Gamma) \cup N'(\Gamma)}$ which maps each i-handle of $\mathfrak{H}$ diffeomorphicically onto the corresponding i-handle of $\mathfrak{H}'$.*

Proof. We use invariance up to isotopy of the tubular neighborhood of a neatly embedded submanifold $Y \subseteq W$, [Hir76, Sec. 6]. We first apply isotopy invariance to each 0-handle, which are tubular neighborhoods of vertices in $V(\Gamma) \cap \operatorname{int} W$. This yields an isotopy Φ_t fixing the vertices which maps each 0-handle of $\mathfrak{H}$ diffeomorphically onto the corresponding 0-handle of $\mathfrak{H}'$. It is easy to see that such an isotopy can be taken to be fixed outside of a neighborhood of the union of the 4-balls associated to each 0-handle. Similarly, it is also easy to arrange the isotopy Φ_t so $\Phi_1 H(e)$ is a regular neighborhood of $e \setminus \cup_{v \in V} \Phi_1 H(v)$. Hence we can assume that the handlebody decompositions $\mathfrak{H}$ and $\mathfrak{H}'$ have 0-handles which are equal. Let $B^0_1, \ldots, B^0_k$ denote the 0-handles of the handlebody decompositions $\mathfrak{H}$ and $\mathfrak{H}'$. Let $e_1, \ldots, e_n$ denote the edges of Γ. We proceed by induction. Suppose that an intermediate isotopy Φ^i_t has been constructed so that all of the 1-handles corresponding to edges e_j with $j \leq i$ have corresponding 1-handles in $\Phi^i_1(\mathfrak{H})$ and $\mathfrak{H}'$ which are equal and such that Φ^i_1 preserves Γ. Let U_i denote the interior of the union of all of the 0-handles and all of the 1-handles corresponding to edges e_j with $j \leq i$ (for either handlebody decomposition $\Phi^i_1(\mathfrak{H})$ or $\mathfrak{H}'$, since these handles agree by induction). Then the 1-handles corresponding to e_{j+1} in the two handelebody decompositions are both regular neighborhoods of a 1-dimensional neatly embedded submanifold of $W \setminus U_i$, and hence there is a (neat) isotopy of $W \setminus U_i$ mapping the 1-handle in $\Phi^i_1(\mathfrak{H})$ to the 1-handle of $\mathfrak{H}'$. It is easy to see that this isotopy can be taken to be fixed outside of a given neighborhood of the union of the two 1-handles corresponding to e_{j+1} in $\Phi^i_1(\mathfrak{H})$ and $\mathfrak{H}'$. The isotopy Φ^{j+1}_t can then be extended to U_i. Proceeding by induction yields the total isotopy as a composition. □

Corollary 11.6. *The graph cobordism maps are invariant under the choice of regular neighborhood $N(\Gamma_0)$ and the choice of a handlebody decomposition of $N(\Gamma_0)$ into 0-handles and 1-handles (the combinatorics of which is determined by the graph Γ_0, as in the construction of the graph cobordism map).*

Proof. One simply uses Lemma 11.5 to create an isotopy which maps one regular neighborhood to the other, then a commutative diagram nearly identical to the one in Lemma 11.4 proves invariance of the graph cobordism map. □

11.5. **Invariance of the graph cobordism maps from subdivision.** One important property of the graph cobordism maps is that they are invariant under subdivision of an edge, with respect to either absolute ordering of the edges adjacent to the new vertex. It follows that the ordering of the edges at a 2-valent vertex does not affect the graph cobordism maps, though more generally we will show in Section 11.6 that the graph cobordism maps are invariant under the cyclic reordering of the edges incident to a vertex.

Lemma 11.7. *Suppose $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ is a graph cobordism with Γ a strong ribbon graph and $\mathfrak{s}\in \mathrm{Spin}^c(W)$, and that v is a point on the interior of Γ. Suppose that Γ' is the graph resulting from adding v to the interior of an edge. Then $F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W,\Gamma',\mathfrak{s}}$. The result holds for either ordering of the two edges adjacent to the new vertex.*

Proof. First note that subdividing an edge does not change the regular neighborhood, so we can take $N(\Gamma_0) = N(\Gamma_0')$. The handlebody decomposition of $N(\Gamma_0)$ does change slightly however. There is an extra 0-handle corresponding to the vertex v, and the edge e which contained v is broken into two edges e_1 and e_2. Let $\mathfrak{H}$ denote the handlebody decomposition of $N(\Gamma_0)$ corresponding to Γ, and let $\mathfrak{H}'$ denote the handlebody decomposition of $N(\Gamma_0') = N(\Gamma_0)$ according to Γ'.

Pick graphs $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}_0'$ inside of $\partial N(\Gamma_0)$ which are neatly isotopic to Γ_0 in $N(\Gamma_0)$ such that $\widetilde{\Gamma}_0'$ is a subdivision of $\widetilde{\Gamma}_0$, as well as a neat isotopy $\Phi_t : W\times I\to W$ mapping Γ_0 to $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}_0'$. Let $\phi = \Phi_1|_{Y_1}$.

Using the formula for the graph cobordism map in Equation (8), we want to show that

$$F^\circ_{\mathcal{W}',\mathfrak{s}|_{W'}}\circ A^V_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}\circ\phi_* = F^\circ_{\mathcal{W}',\mathfrak{s}|_{W'}}\circ A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'}\circ F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}\circ\phi_*.$$

To this end, it is sufficient to show that

$$A^V_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} = A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'}\circ F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}.$$

Applying Lemma 7.10, we see that for either ordering of the edges e_1 and e_2, the desired equality becomes

$$A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'}\circ S^+_v\circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} = A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'}\circ F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}.$$

To show this, it is sufficient to show that

$$S^+_v\circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} = F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}.$$

This is a relatively easy model computation. In Figure 17, we perform a model computation of $F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}$ and see explicitly that this is equal to $S^+_v\circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}$, completing the proof.

□

11.6. **Dependence on cyclic ordering.** In the construction of the graph cobordism maps, we picked an absolute ordering of the edges incident to a vertex representing a given cyclic ordering. In this section, we show that the graph cobordism maps are invariant under cyclic reordering of the edges adjacent to a given vertex. The graph cobordism map is defined as a composition of several maps, including a graph action map. Although the graph action map does depend on the absolute ordering since we have included all of the vertices in the incoming set V, the composition does not:

Lemma 11.8. *The graph cobordism maps depend only on the cyclic ordering of the edges incident to a vertex.*

Proof. Let $\Gamma_0, \widetilde{\Gamma}_0$ and $\mathfrak{H}$ be as in the construction of the graph cobordism map. We will show that $A_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H}}$ depends only on the cyclic ordering at a vertex. To do this, we will factor $A_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H}}$ in an interesting way, and then apply Lemma 7.12.

Let $\mathcal{G}$ denote the flow graph that we associate to $\widetilde{\Gamma}_0$ in the graph cobordism maps and let $\mathcal{G}_m\circ\cdots\circ\mathcal{G}_1$ denote a Cerf decomposition of $\mathcal{G}$. Let $\mathcal{G}_v$ denote the elementary flow graph associated to the vertex v. Pick two distinct edges e_i and e_j which are incident to v, arbitrarily. By Lemma 11.7, subdividing the graph Γ does not effect the graph cobordism map, so we can assume that e_i and e_j are distinct edges which both connect v to vertices which are not v. Let $\mathfrak{H}'$ denote the handlebody decomposition of $N(\Gamma_0)$ obtained by replacing the union of the 0-handle associated to v and the 1-handles associated to e_i and e_j, with a single 1-handle. We now claim that

$$F^\circ_{\mathfrak{H}} = S^+_v\circ F_{\mathfrak{H}'}.$$

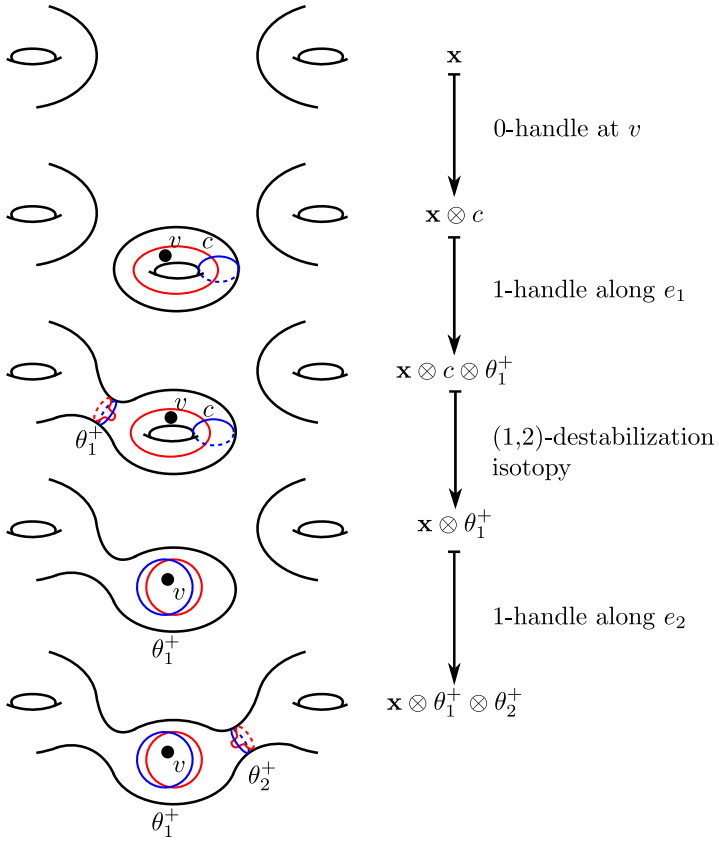

FIGURE 17. A computation of a portion of the maps in $F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}$ in Lemma 11.7. We see explicitly that the composition is $S_v^+ \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}$.

This is essentially the same model computation that we used to prove subdivision invariance in Lemma 11.7, and is shown in Figure 17. Having established the above equality, letting $\widetilde{e}_i$ denote the portion of the edge e_i in the elementary flow graph $\mathcal{G}_v$, we now just compute that

$$\begin{aligned} A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H}} &= A_{\mathcal{G}_m} \cdots A_{\mathcal{G}_v} \cdots A_{\mathcal{G}_1} \circ F^\circ_{\mathfrak{H}} \\ &= A_{\mathcal{G}_m} \cdots (S_v^- A_{\widetilde{e}_n} \cdots A_{\widetilde{e}_1}) \cdots A_{\mathcal{G}_1} \circ S_v^+ \circ F^\circ_{\mathfrak{H}'} \\ &= A_{\mathcal{G}_m} \cdots (S_v^- A_{\widetilde{e}_n} \cdots A_{\widetilde{e}_1} S_v^+) \cdots A_{\mathcal{G}_1} \circ F^\circ_{\mathfrak{H}'}. \end{aligned}$$

The only term in the above expression that depends on the ordering of the edges $e_1, \ldots, e_n$ is $(S_v^- A_{\widetilde{e}_n} \cdots A_{\widetilde{e}_1} S_v^+)$, so applying Lemma 7.12, shows that it only depends on the cyclic ordering of the edges. □

12. The composition law for $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$

The main goal of this section is to prove Lemma 12.7, i.e. the composition law for graph cobordisms $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ with $\pi_0(Y_i) \to \pi_0(W)$ a surjection for $i = 1, 2$. In working towards a proof of this, much of this section is devoted to showing the relationship between the graph cobordism maps, the graph action maps and the path cobordism maps. The simplest such result is Lemma 12.3, which states that the graph cobordism maps for a cylindrical cobordism $(Y \times I, \Gamma)$ are the same as a the graph action map for the graph projected into Y. We begin with some technical lemmas:

Lemma 12.1. *Suppose that $(W, \boldsymbol{\gamma}) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a path cobordism and E is a neatly embedded path in W between two points in Y_1. Let $N(E)$ denote a tubular neighborhood of E and let $\mathfrak{E}$ denote the handlebody decomposition of $N(E)$ as a single 1-handle as in the construction of the graph cobordism maps.*

Let $\mathfrak{s} \in \mathrm{Spin}^c(W)$. Letting $F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}$ denote the handlebody map in the construction of the graph cobordism maps, then

$$F^\circ_{W\setminus N(E),\boldsymbol{\gamma},\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}}.$$

Proof. By construction

$$F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} = F_{(Y_1\times I)\cup N(E),\mathbf{w}_1\times I,\mathfrak{s}'},$$

where $\mathfrak{s}'$ is the unique Spin^c structure on $Y_1 \times I \cup N(E)$ extending $\mathfrak{s}|_{Y_1}$. By the composition law for path cobordism maps (Corollary 10.3) we have

$$F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} = F^\circ_{W\cup(Y_1\times I),\boldsymbol{\gamma}\cup(\mathbf{w}_1\times I),\overline{\mathfrak{s}}},$$

where $\overline{\mathfrak{s}}$ is the unique Spin^c structure on $W \cup (Y_1 \times I)$ extending $\mathfrak{s}$. Using this and the composition law for path cobordisms, we have that

$$\begin{aligned} F_{W,\boldsymbol{\gamma},\mathfrak{s}} &= F^\circ_{W\cup(Y_1\times I),\boldsymbol{\gamma}\cup(\mathbf{w}_1\times I),\overline{\mathfrak{s}}} \\ &= F^\circ_{W\setminus N(E),\boldsymbol{\gamma},\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{(Y_1\times I)\cup N(E),(\mathbf{w}_1\times I),\overline{\mathfrak{s}}} \\ &= F^\circ_{W\setminus N(E),\boldsymbol{\gamma},\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}, \end{aligned}$$

as we wanted. □

Lemma 12.2. *Suppose that W is a cobordism from Y_1 to Y_2 with a Morse function f which has critical points which are only index 1,2 or 3 and such that any index i critical point has value less than any index $i+1$ critical point. Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \varnothing)$ is a flow graph in Y_1, which has trivial intersection with the descending manifolds of the critical points of f. Let $\boldsymbol{\gamma}_0$ denote the paths determined by the flow under f of the points V_0 and similarly let $\boldsymbol{\gamma}_1$ denote the paths determined by the flow of the points V_1. Then there is a flow graph $\mathcal{G}' = (\Gamma', V_0', V_1', \sigma, \varnothing)$ in Y_2 such that*

$$F^\circ_{W,\boldsymbol{\gamma}_1,\mathfrak{s}} \circ A_{\mathcal{G}} = A_{\mathcal{G}'} \circ F^\circ_{W,\boldsymbol{\gamma}_0,\mathfrak{s}}.$$

After perturbing to be neatly embedded in W, $\Gamma' * \boldsymbol{\gamma}_0$ and $\boldsymbol{\gamma}_1 * \Gamma$ are isotopic.

Proof. Using the composition law for path cobordisms we can assume that f either has a single critical point of index 1 or 3, or has only critical points of index 2. The Morse function f induces a diffeomorphism $\psi_f : Y_1(\mathbb{L}) \to Y_2$, where $\mathbb{L}$ is the collection of framed spheres (all of the same dimension) consisting of the the intersection of the descending manifolds of critical points of f and Y_1. Assuming that Γ intersects these framed spheres trivially, the map ψ_f then naturally induces a flow graph $\mathcal{G}' = (\psi_f)_*\mathcal{G}$, which is the image of $\mathcal{G}$ under ψ_f. The paths $\boldsymbol{\gamma}_i$ are simply the flow under the gradient flow of f of the basepoints V_i. The basepoints V_i' are the images under ψ_f of V_i.

To see that we have the relation

$$F^\circ_{W,\boldsymbol{\gamma}_1,\mathfrak{s}} \circ A_{\mathcal{G}} = A_{\mathcal{G}'} \circ F^\circ_{W,\boldsymbol{\gamma}_0,\mathfrak{s}},$$

we pick a Heegaard surface $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, V_0)$ for (Y_1, V_0). We can isotope Σ so that V_1 sits inside of Σ. Now by stabilizing the Heegaard surface Σ and isotoping the edges and additional vertices of Γ relative the points V_0 and V_1, we can assume they sit inside of Σ. We can assume that the final image of the graph does intersect the embedded spheres corresponding to the descending manifold of the critical points of f, since we assumed that the graph Γ initially did not intersect those spheres. The Morse function f also induces a diffeomorphism of $\Phi : W \to W(Y_1, \mathbb{L})$. Given a nonempty collection of basepoints $\mathbf{w} \subseteq \Sigma$, if $\boldsymbol{\gamma}_{\mathbf{w}}$ is the collection of paths in W given by flowing $\mathbf{w}$ along the gradient flow of f, then by definition

$$F^\circ_{W,\boldsymbol{\gamma}_{\mathbf{w}},\mathfrak{s}} = \psi_f \circ F^\circ_{\mathbb{L},\Phi_*\mathfrak{s}}.$$

In particular, since the maps associated to framed 0-spheres, framed links, and framed 2-spheres all commute with the free stabilization and destabilization maps (Theorem 6.17, Lemma 8.7) , as well as the relative

homology maps (Lemmas 5.4 and 8.6), we compute that

$$\begin{aligned} F^{\circ}_{W,\boldsymbol{\gamma}_1,\mathfrak{s}} \circ A_{\mathcal{G}} &= \psi_f \circ F^{\circ}_{\mathbb{L},\Phi_*\mathfrak{s}} \circ A_{\mathcal{G}} \\ &= \psi_f \circ A_{\mathcal{G}} \circ F^{\circ}_{\mathbb{L},\Phi_*\mathfrak{s}} \\ &= A_{\mathcal{G}'} \circ \psi_f \circ F^{\circ}_{\mathbb{L},\Phi_*\mathfrak{s}} \\ &= A_{\mathcal{G}'} \circ F^{\circ}_{W,\boldsymbol{\gamma}_0,\mathfrak{s}}, \end{aligned}$$

as we wanted to show.

As graphs which are perturbed to be neatly embedded in W, $\Gamma' * \boldsymbol{\gamma}_0$ and $\boldsymbol{\gamma}_1 * \Gamma$ can be seen to be isotopic in W because Γ' is simply the flow of Γ under the gradient flow of f. □

Lemma 12.3. *Suppose that $(W,\Gamma) : (Y,\mathbf{w}_1) \to (Y,\mathbf{w}_2)$ is a graph cobordism with $W = Y \times I$. Assume that Γ is perturbed so that Γ projects to an embedded graph Γ_Y in Y under the projection $\pi : Y \times I \to Y$. Then*

$$F^{\circ}_{W,\Gamma,\mathfrak{s}} = A_{\mathcal{G},\mathfrak{s}}$$

where $\mathcal{G} = (\Gamma_Y, \mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$.

Proof. Let $\Gamma_0, N(\Gamma_0), \widetilde{\Gamma}_0, \phi_*, \widetilde{\boldsymbol{\gamma}}_f$ and $\mathfrak{H}$ be as in the construction of the graph cobordism maps. Let

$$V_W = V(\Gamma_0) \cap \operatorname{int} W$$

denote the collection of vertices of Γ_0 which are in the interior of W. Let V_Y denote the projection of these vertices onto $Y \times \{0\}$. Let $\widetilde{V}_W$ denote the vertices of $\widetilde{\Gamma}_0$ corresponding to vertices of V_W. By assuming Γ is in generic position, we can pick $\widetilde{v} \in \widetilde{V}_W$ so that $\pi_Y \widetilde{v} = \pi_Y v$. Pick a set of paths E in W from points in $Y \times \{0\} \setminus N(\mathbf{w})$ to the vertices in V_W such that the paths $e \in E$ are constant in the Y component. By picking an appropriate regular neighborhood $N(\Gamma_0)$, we can assume that the paths E naturally induce a set of paths E' in $W \setminus N(\Gamma_0)$ from points in $Y \times \{0\} \setminus N(\mathbf{w}_1)$ to points in $\partial N(v)$ for each $v \in V_W$. Let $\mathfrak{E}$ denote the handle decomposition of $N(E')$ induced by these paths.

Pick $\widetilde{\Gamma}_0$ and $\widetilde{\boldsymbol{\gamma}}_f$ so that $\widetilde{\boldsymbol{\gamma}}_f$ are also constant in the Y component. Let $\mathbf{w}_f$ denote the basepoints in $\widetilde{\boldsymbol{\gamma}}_f \cap \widetilde{\Gamma}_0$. Pick paths $\boldsymbol{\gamma}'$ from the vertices $\widetilde{v} \in \widetilde{V}_W \setminus \mathbf{w}_f$ to points in $Y \times \{1\}$ and pick the paths in $\boldsymbol{\gamma}'$ to all be constant in Y as well.

By definition we have that

$$F^{\circ}_{W,\Gamma,\mathfrak{s}} = F_{W \setminus N(\Gamma_0), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}|_{W \setminus N(\Gamma_0)}} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}} \circ \phi_*.$$

Write $\mathfrak{H} = \mathfrak{H}_0 \cup \mathfrak{H}_1$ where $\mathfrak{H}_0$ are the 0-handles corresponding to vertices in V_W and $\mathfrak{H}_1$ are the 1-handles corresponding to edges of Γ_0. Using the above observation, and Lemma 12.1, we perform the following manipulation (in the following, we will abbreviate any restriction of $\mathfrak{s}$ by $\mathfrak{s}|$):

$$\begin{aligned} F_{W \setminus N(\Gamma_0), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}} \circ \phi_* &= F^{\circ}_{W \setminus (N(\Gamma_0) \cup N(E)), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}|} \circ F^{\circ}_{\mathfrak{E}} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ F^{\circ}_{\mathfrak{H}_0} \circ \phi_* \\ &= F^{\circ}_{W \setminus (N(\Gamma_0) \cup N(E)), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ F^{\circ}_{\mathfrak{E}} \circ F^{\circ}_{\mathfrak{H}_0} \circ \phi_*. \end{aligned}$$

The second manipulation results from commuting $F^{\circ}_{\mathfrak{E}}$ with the graph action map and the 1-handle maps in $F^{\circ}_{\mathfrak{H}_1}$. Letting $\widetilde{\mathbf{w}}_1 = \phi(\mathbf{w}_1)$, there is a diffeomorphism $\Phi : (Y, \widetilde{\mathbf{w}}_1 \cup V_Y) \to (Y \# \partial(N(E) \cup N(V_W)), \widetilde{\mathbf{w}}_1 \cup \widetilde{V}_W)$ mapping a vertex $\widetilde{v} \in V_Y$ to the corresponding $v \in \widetilde{V}_W$. The map Φ can be taken to be the identity outside a small neighborhood of $(Y \# \partial(N(E) \cup N(V_W))) \setminus Y$. A model computation similar to the one in Lemma 11.7 shows that

$$F^{\circ}_{\mathfrak{E}} \circ F^{\circ}_{\mathfrak{H}_0} = \Phi_* \circ S^{+}_{V_Y}.$$

Continuing the manipulations above, we thus have that

$$F^{\circ}_{W \setminus (N(\Gamma_0) \cup N(E)), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ F^{\circ}_{\mathfrak{E}} \circ F^{\circ}_{\mathfrak{H}_0} \circ \phi_* = F^{\circ}_{W \setminus (N(\Gamma_0) \cup N(E)), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ S^{+}_{V_Y} \circ \phi_*.$$

We now claim that $W \setminus (N(\Gamma_0) \cup N(E))$ has Morse function which has only index 2 critical points, and which has the paths in $\widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}'$ as flow lines. To see this, first observe that there is a diffeomorphism of $W \setminus (N(E) \cup N(V_W))$ with $W = Y \times I$ which extends Φ and maps the paths $\widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}'$ to paths of the form $(V_Y \cup \widetilde{\mathbf{w}}_1) \times I$. On $W = Y \times I$, we first take the Morse function $f_0(y,t) = t$. We then modify f_0 to create a Morse function f by introducing an index 1-2 critical point pair for each edge of Γ_0. We then compose f_0

with an isotopy of W which moves the descending manifold of each index 1 critical point to the core of the corresponding 1-handle in $\mathfrak{H}_1$. This can be done without affecting the flowlines corresponding to the points $V_Y \cup \widetilde{\mathbf{w}}_1$. Such a Morse function f_0 induces a decomposition $W = W_2 \cup_{Y'} W_1$ where Y' is a level set of f_0 and $f_0|_{W_1}$ has only index 1 critical points and $f_0|_{W_2}$ has only index 2 critical points. Furthermore, there is an equivalence

$$\Psi : (W \setminus (N(\Gamma_0) \cup N(E)), \widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}') \to (W_2, (\widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}') \cap W_2).$$

We then define a Morse function f on $W \setminus (N(\Gamma_0) \cup N(E))$ by $f = f_0 \circ \Psi$. Let ψ_f denote the diffeomorphism

$$\psi_f : (Y \# \partial(N(E) \cup N(\Gamma_0))(\mathbb{L}), \widetilde{\mathbf{w}}_1 \cup \widetilde{V}_W) \to (Y, \widetilde{\mathbf{w}}_1 \cup V_Y)$$

induced by the Morse function $f|_{W_2}$, where $\mathbb{L} \subseteq Y \# \partial(N(E) \cup N(\Gamma_0))$ is a framed link.

We now observe that if $\mathbb{L}_0 \subseteq Y_0$ is a framed link in a 3-manifold Y_0, and $\mathcal{G} = (\Gamma, V_0, V_1, V, \sigma)$ is a flow graph in Y_0, then

$$F^\circ_{\mathbb{L}_0,\mathfrak{s}} \circ A_{\mathcal{G}} = A_{\mathcal{G}'} \circ F^\circ_{\mathbb{L}_0,\mathfrak{s}}$$

as maps from $HF^\circ(Y_0, V, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_0(\mathbb{L}), V_1, \mathfrak{s}|_{Y_0(\mathbb{L})})$, where $\mathcal{G}'$ is the induced map on $Y_0(\mathbb{L}_0)$. This follows since the free stabilization maps and the relative homology maps commute with the triangle maps (Theorem 6.17 and Lemma 5.4).

We thus have that

$$\begin{aligned}
F^\circ_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* &= \psi_f \circ F^\circ_{\mathbb{L},\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* \\
&= \psi_f \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathbb{L},\mathfrak{s}|} \circ F^\circ_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ \psi_f \circ F^\circ_{\mathbb{L},\mathfrak{s}|} \circ F^\circ_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ \psi_f \circ F^\circ_{\mathbb{L},\mathfrak{s}|} \circ F^\circ_{\mathfrak{H}_1} \circ \Phi_* \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^\circ_{W\setminus(N(\Gamma_0)\cup N(E),((\Phi\circ\phi)_*\mathbf{w}_1)\times I,\mathfrak{s}|} \\
&\quad \circ F^\circ_{\mathfrak{H}_1} \circ \Phi_* \circ \phi_*.
\end{aligned}$$

By Lemma 12.1, we have

$$A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^\circ_{W\setminus(N(\Gamma_0)\cup N(E),(\Phi\circ\phi)_*\mathbf{w}_1,\mathfrak{s}|} \circ F^\circ_{\mathfrak{H}_1} \circ \Phi_* \circ \phi_* = A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^\circ_{W\setminus N(E),(\Phi\circ\phi)_*\mathbf{w}_1,\mathfrak{s}|} \circ \Phi_* \circ \phi_*$$

Since $\Phi : Y \to Y \# \partial N(E)$ extends to an equivalence of $W \setminus N(E)$ with $W = Y \times I$ which is the identity on $Y \times \{1\}$, we have

$$\begin{aligned}
A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^\circ_{W\setminus N(E),(\Phi\circ\phi)_*\mathbf{w}_1\times I,\mathfrak{s}|} \circ \Phi_* \circ \phi_* &= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^\circ_{W,\phi_*\mathbf{w}_1\times I,\mathfrak{s}|} \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ \phi_*.
\end{aligned}$$

At this point we are almost done. Note that in the above expression $A_{(\psi_f)_*\widetilde{\Gamma}_0}$ is the graph action map corresponding to a flow graph $\mathcal{G}'_f = ((\psi_f)_*\widetilde{\Gamma}_0, V_Y \cup \phi_*\mathbf{w}_1, \phi_*\mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$. Letting $\mathcal{G}_f = ((\psi_f)_*\widetilde{\Gamma}_0, \phi_*\mathbf{w}_1, \phi_*\mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$, by Lemma 7.10, we know that

$$\begin{aligned}
A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ \phi_* &= A_{\mathcal{G}'_f} \circ S^+_{\psi_f V_Y} \circ \phi_* \\
&= A_{\mathcal{G}_f} \circ \phi_*.
\end{aligned}$$

It remains to see what the graph $(\psi_f)_*\widetilde{\Gamma}_0$ is. Note that the diffeomorphism ψ_f is induced by the flow of f. The flow of f preserves the basepoints $V_Y \cup \phi_*\mathbf{w}_1$ and the images of the edges of $\widetilde{\Gamma}_0$ are induced by the flow of f. Hence if e is an edge of Γ and e_Y is its projection, then e_Y, e and $\psi_f(e)$ are all isotopic via an isotopy which maps the endpoints of e into $(\phi_*\mathbf{w}_1 \cup V_Y) \times I$. In particular the relative homology map in Y along the edge e_Y and $\psi_f(e)$ agree. If $\mathcal{G} = (\Gamma, \mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$, we thus have that

$$A_{\mathcal{G}_f} = A_{\phi_*\mathcal{G}},$$

and hence combining all of the computations we've done so far, we thus have that

$$\begin{aligned} F^\circ_{W,\Gamma,\mathfrak{s}} &= A_{\phi_*\mathcal{G}} \circ \phi_* \\ &= \phi_* \circ A_{\mathcal{G}} \\ &= A_{\mathcal{G}}. \end{aligned}$$

The last equality follows since ϕ is isotopic to the identity relative the basepoints $\mathbf{w}_2$. □

Lemma 12.4. *Suppose that $(W,\Gamma') = (W_2,\boldsymbol{\gamma}) \circ (W_1,\Gamma)$ where $(W_2,\boldsymbol{\gamma}_2)$ is a path cobordism. Then*

$$F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}} F^\circ_{W,\Gamma',\mathfrak{s}},$$

where $F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2}$ is the path cobordism map, and the remaining maps in the above expression are graph cobordism maps.

Proof. This follows essentially from the construction of the graph cobordism maps and the composition law for path cobordisms (Corollary 10.3). Pick Γ_0, $N(\Gamma_0), \widetilde{\Gamma}_0, \phi_*, \widetilde{\boldsymbol{\gamma}}_f$ and $\mathfrak{H}$ as in the construction of the graph cobordism maps. Then by construction

$$F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma,\mathfrak{s}_1} = F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F_{W_1\setminus N(\Gamma_0),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}_1|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}_1|_{Y_1}} \circ \phi_*.$$

Writing $(W',\boldsymbol{\gamma}')$ for $(W_2,\boldsymbol{\gamma}) \circ (W_1 \setminus N(\Gamma_0), \widetilde{\boldsymbol{\gamma}}_f)$ we see that

$$F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F_{W_1\setminus N(\Gamma_0),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}_1|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}_1|_{Y_1}} \circ \phi_*$$

$$\left(\sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W') \\ \mathfrak{s}|_{W_2}=\mathfrak{s}_2 \\ \mathfrak{s}|_{W_1\setminus N(\Gamma)_0}=\mathfrak{s}_1|_{W_1\setminus N(\Gamma_0)}}} F_{W',\boldsymbol{\gamma}',\mathfrak{s}} \right) \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}_1|_{Y_1}} \circ \phi_*,$$

but any Spin^c structure on W' extends uniquely over $N(\Gamma_0)$ to a Spin^c structure on W, since W is obtained by adding 3- and 4-handles to W'. Hence the above expression is equal to

$$\left(\sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_2}=\mathfrak{s}_2 \\ \mathfrak{s}|_{W_1}=\mathfrak{s}_1}} F^\circ_{W',\boldsymbol{\gamma}',\mathfrak{s}|_{W'}} \right) \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} \circ \phi_*,$$

which by definition is equal to

$$\sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_2}=\mathfrak{s}_2 \\ \mathfrak{s}|_{W_1}=\mathfrak{s}_1}} F^\circ_{W,\Gamma',\mathfrak{s}},$$

as claimed. □

Lemma 12.5. *Suppose that $(W,\Gamma) : (Y_1,\mathbf{w}_1) \to (Y_2,\mathbf{w}_2)$ is a graph cobordism and suppose that Γ is isotopic (necessarily not neatly) to a graph of the form $\Gamma' \cup \boldsymbol{\gamma}$ where $\Gamma' \subseteq Y_1$ is a flow graph and $\boldsymbol{\gamma}$ is a collection of paths from the terminal basepoints of Γ' in Y_1 to the basepoints $\mathbf{w}_2 \subseteq Y_2$. Then*

$$F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} \circ A_{\Gamma'}.$$

Proof. Pick a collar neighborhood $N(Y_1) = Y_1 \times I$ inside of W. By assumption we can isotope Γ so that W is equivalent to the composition $(W,\boldsymbol{\gamma}) \circ (N(Y_1),\Gamma'')$ where Γ'' is a graph in $N(Y_1) \times I$ which projects to Γ'. By Lemmas 12.4 and 12.3 we know that

$$F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} \circ F^\circ_{N(Y_1),\Gamma'',\mathfrak{s}|_{N(Y_1)}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} \circ A_{\Gamma'},$$

as we wanted. □

Lemma 12.6. *Suppose $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ is a graph cobordism with $\pi_0(Y_i)\to\pi_0(W)$ surjective. Then there is a decomposition of W into*

$$(W,\Gamma)=(W_3,\boldsymbol{\gamma}_3)\circ(W_2,\boldsymbol{\gamma}_2)\circ(N,\Gamma')\circ(W_1,\boldsymbol{\gamma}_2),$$

where W_i has a Morse function with only index-i critical points and $N\cong Y'\times I$ is a cylindrical cobordism. Furthermore

$$F^\circ_{W,\Gamma,\mathfrak{s}}=F^\circ_{W_3,\boldsymbol{\gamma}_3,\mathfrak{s}|_{W_3}}\circ F^\circ_{W_2,\boldsymbol{\gamma}_2,\mathfrak{s}|_{W_2}}\circ A_{\Gamma'_{Y'}}\circ F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}}.$$

Proof. First pick a Morse function f on W with critical points of increasing index and no index 0 or index 4 critical points, which is possible since $\pi_0(Y_i)\to\pi_0(W)$ is surjective for $i=1,2$. By picking appropriate level sets of f, we get a decomposition of W into cobordisms $W=W_3\cup W_2\cup N\cup W_1$ where $f|_{W_i}$ is a Morse function with only index i critical points and $f|_N$ has no critical points. Writing Y' for the common boundary between W_1 and N, we first note that $\pi_1(Y',q)\to\pi_1(W,q)$ is surjective for any choice of q. As a consequence, we can isotope Γ within W and perturb f slightly so that $\Gamma\cap W_i$ is a collection of paths which are flowlines of f.

Let W' denote $W_3\cup W_2\cup N$ and let (W',Γ') denote the graph cobordism $(W',\Gamma\cap W')$. Let the basepoints of $Y'=N\cap W_1$ be denoted by $\mathbf{w}'$. By Lemmas 12.4, 12.5 we see that

$$F^\circ_{W',\Gamma',\mathfrak{s}'}=F^\circ_{W_3,\boldsymbol{\gamma}_3,\mathfrak{s}|_{W_3}}\circ F^\circ_{W_2,\boldsymbol{\gamma}_2,\mathfrak{s}|_{W_2}}\circ A_{\Gamma'}.$$

Hence it is sufficient to show that

$$F^\circ_{W',\Gamma',\mathfrak{s}|_{W'}}\circ F_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}}=F^\circ_{W,\Gamma,\mathfrak{s}}.$$

For this, let E consist of all of the descending manifolds of the index 1 critical points of f inside of W. Let $\mathfrak{E}$ denote the induced handle decomposition of $N(E)$ with one 1-handle for each arc of E. The flow of f yields a diffeomorphism of $\Phi:(Y_1\#\partial N(E),\mathbf{w}_1)\to(Y',\mathbf{w}')$. As in Lemma 12.1 the 1-handle map $F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}}$ satisfies

$$F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}}=\Phi_*\circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}.$$

On the other hand, there is an equivalence

$$\widehat{\Phi}:(W\setminus N(E),\Gamma)\to(W',\Gamma')$$

extending Φ. Hence we have

$$\begin{aligned}F^\circ_{W',\Gamma',\mathfrak{s}'}\circ F_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}}&=F^\circ_{W',\Gamma',\mathfrak{s}'}\circ\Phi_*\circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}\\&=F^\circ_{W\setminus N(E),\Gamma,\mathfrak{s}|_{W\setminus N(E)}}\circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}\end{aligned}$$

Let $\Gamma_0,\widetilde{\Gamma}_0,N(\Gamma_0),\phi,\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{H}$ be chosen as in the construction of the graph cobordism maps for the graph cobordism $(W\setminus N(E),\Gamma)$. By construction

$$F^\circ_{W\setminus N(E),\Gamma,\mathfrak{s}|_{W\setminus N(E)}}=F^\circ_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|}\circ A_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H},\mathfrak{s}|}\circ\phi_*.$$

Hence we have that

$$\begin{aligned}F^\circ_{W\setminus N(E),\Gamma,\mathfrak{s}|_{W\setminus N(E)}}\circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}&=F^\circ_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|}\circ A_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H},\mathfrak{s}|}\circ\phi_*\circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}\\&=F^\circ_{W\setminus(N(E)\cup N(\Gamma_0)),\Gamma,\mathfrak{s}|_{W\setminus N(E)}}\circ F^\circ_{\mathfrak{E},\mathfrak{s}|}\circ A_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H},\mathfrak{s}|}\circ\phi_*.\end{aligned}$$

The last equality follows since the map $F^\circ_{\mathfrak{E},\mathfrak{s}|}$ commutes with 0-handle, 1-handle maps, the free stabilization maps, and the relative homology maps. It also commutes with ϕ_* (or more precisely a choice of two different ϕ maps, one on Y_1 and one on $Y_1\#\partial N(E)$) because ϕ_* is fixed outside of a neighborhood of the basepoints, and the 1-handles in $\mathfrak{E}$ are attached away from the basepoints. Finally, by applying Lemma 12.1, we see that

$$\begin{aligned}F^\circ_{W\setminus(N(E)\cup N(\Gamma_0)),\Gamma,\mathfrak{s}|}\circ F^\circ_{\mathfrak{E},\mathfrak{s}|}\circ A_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H},\mathfrak{s}|}\circ\phi_*&=F^\circ_{W\setminus N(\Gamma_0)),\Gamma,\mathfrak{s}|}\circ A_{\widetilde{\Gamma}_0}\circ F^\circ_{\mathfrak{H},\mathfrak{s}|}\circ\phi_*\\&=F^\circ_{W,\Gamma,\mathfrak{s}},\end{aligned}$$

completing the proof. □

Lemma 12.7. *Suppose that $(W,\Gamma) = (W_2,\Gamma_2)\circ(W_1,\Gamma_1)$ is a decomposotion of a graph cobordism. Writing $(W_i,\Gamma_i):(Y_i,\mathbf{w}_i)\to(Y_i',\mathbf{w}_i)$, suppose that $\pi_0(Y_i)\to\pi_0(W_i)$ and $\pi_0(Y_i')\to\pi_0(W_i)$ are all surjective. Then*

$$F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2}\circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in\mathrm{Spin}^c(W)\\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}} F^\circ_{W,\Gamma,\mathfrak{s}}.$$

Proof. Using Lemma 12.6 to decompose each cobordism W_i into a composition

$$(W_{3i},\boldsymbol{\gamma}_{3i})\circ(W_{2i},\boldsymbol{\gamma}_{2i})\circ(N_i,\Gamma_i')\circ(W_{1i},\boldsymbol{\gamma}_{1i}),$$

by the same Lemma we have that

$$F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2}\circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1}$$
$$= F^\circ_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}}\circ F^\circ_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}}\circ A_{\Gamma'_{2,\pi}}\circ F^\circ_{W_{12},\boldsymbol{\gamma}_{12},\mathfrak{s}_{12}}\circ F^\circ_{W_{31},\boldsymbol{\gamma}_{31},\mathfrak{s}_{31}}\circ F^\circ_{W_{21},\boldsymbol{\gamma}_{21},\mathfrak{s}_{21}|}\circ A_{\Gamma'_{1,\pi}}\circ F^\circ_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}},$$

where $\mathfrak{s}_{ij} = \mathfrak{s}_j|_{W_{ij}}$ and $A_{\Gamma'_{i,\pi}}$ denotes the projection of Γ_i' into the the incoming end of N_i given by the flow of the Morse function f.

We can assume that such a decomposition resulted from the level sets of a Morse function f_i on W_i, and assume further that f_i are the restrictions of a Morse function on $W = W_2\cup W_1$. Also we can assume that $\mathfrak{s}_{12}|_{Y_2} = \mathfrak{s}_{31}|_{Y_1'}$ since otherwise the composition is zero, which satisfies the composition law.

The Morse function $f|_{W_{12}\cup W_{31}}$ has index 1 and index 3 critical points. By using a standard rearrangement theorem we can replace $f|_{W_{21}\cup W_{31}}$ with a Morse function f' which has only index 1 and index 3 critical points, but for which the index 1 critical points have lower value than the index 3 critical points. We can assume that f' still has $\boldsymbol{\gamma}_{12}*\boldsymbol{\gamma}_{31}$ as flowlines since the rearrangement can be assumed to modify the Morse function $f|_{W_{1,2}\cup W_{3,1}}$ only in a neighborhood of the ascending and descending manifolds of the critical points. This allows us to rewrite $(W_{12}\cup W_{31},\boldsymbol{\gamma}_{12}*\boldsymbol{\gamma}_{31})$ as $(W'_{31}\cup W'_{12},\boldsymbol{\gamma}'_{31}*\boldsymbol{\gamma}'_{12})$ where W'_{31} has a Morse function with only index 3 critical points and W'_{12} has a Morse function with only index 1 critical points.

Using the composition rule for path cobordisms (Corollary 10.3) for the cobordism $W_{12}\cup W_{31} = W'_{31}\cup W'_{12}$, we get that

$$F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2}\circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1}$$
$$= F^\circ_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}}\circ F^\circ_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}}\circ A_{\Gamma'_{2,\pi}}\circ F^\circ_{W'_{31},\boldsymbol{\gamma}'_{31},\mathfrak{s}'_{31}}\circ F^\circ_{W'_{12},\boldsymbol{\gamma}'_{12},\mathfrak{s}'_{12}}\circ F^\circ_{W_{21},\boldsymbol{\gamma}_{21},\mathfrak{s}_{21}}\circ A_{\Gamma'_{1,\pi}}\circ F^\circ_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}.$$

Similarly commuting the 1-handle cobordism $(W_{12},\boldsymbol{\gamma}_{12})$ across W_{21} we get that the above expression is equal to

$$F^\circ_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}}\circ F^\circ_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}}\circ A_{\Gamma'_{2,\pi}}\circ F^\circ_{W'_{31},\boldsymbol{\gamma}'_{31},\mathfrak{s}'_{31}}\circ F^\circ_{W'_{21},\boldsymbol{\gamma}'_{21},\mathfrak{s}'_{21}}\circ F^\circ_{W''_{12},\boldsymbol{\gamma}''_{12},\mathfrak{s}''_{12}}\circ A_{\Gamma'_{1,\pi}}\circ F^\circ_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}$$

for cobordisms $(W''_{12},\boldsymbol{\gamma}''_{12})$ and $(W'_{21},\boldsymbol{\gamma}'_{21})$. Applying Lemma 12.2, we commute the graph action maps across the cobordism maps to see that the above expression is equal to

$$F^\circ_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}}\circ F^\circ_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}}\circ F^\circ_{W'_{31},\boldsymbol{\gamma}''_{31},\mathfrak{s}'_{31}}\circ F^\circ_{W'_{21},\boldsymbol{\gamma}''_{21},\mathfrak{s}'_{21}}\circ A_{\Gamma''_{2,\pi}}\circ A_{\Gamma''_{1,\pi}}\circ F^\circ_{W''_{12},\boldsymbol{\gamma}''_{12},\mathfrak{s}''_{12}}\circ F^\circ_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}$$

for paths $\boldsymbol{\gamma}''_{31},\boldsymbol{\gamma}''_{21}$ and $\boldsymbol{\gamma}''_{12}$ and graphs $\Gamma''_{2,\pi}$ and $\Gamma''_{1,\pi}$. By perturbing the Morse functions, graphs and paths slightly we may assume that $\Gamma''_{1,\pi}*\Gamma''_{2,\pi}$ is an embedded flow graph, where $*$ denotes concatenation at initial and terminal vertices. By construction the graph action map is functorial under concatenation, and hence $A_{\Gamma''_{2,\pi}}\circ A_{\Gamma''_{1,\pi}} = A_{\Gamma''_{2,\pi}*\Gamma''_{1,\pi}}$. We can also perform a critical point switch between $(W_{22},\boldsymbol{\gamma}_{22})$ and $(W'_{31},\boldsymbol{\gamma}''_{31})$ to decompose $(W_{22},\boldsymbol{\gamma}_{22})\circ(W'_{31},\boldsymbol{\gamma}''_{31})$ as $(W''_{31},\boldsymbol{\gamma}'''_{31})\circ(W''_{22},\boldsymbol{\gamma}_{22})$. Hence the above expression becomes

$$F^\circ_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}}\circ F^\circ_{W''_{31},\boldsymbol{\gamma}'''_{31},\mathfrak{s}''_{31}}\circ F^\circ_{W'_{22},\boldsymbol{\gamma}'_{22},\mathfrak{s}'_{22}}\circ F^\circ_{W'_{21},\boldsymbol{\gamma}''_{21},\mathfrak{s}'_{21}}\circ A_{\Gamma''_{2,\pi}*\Gamma''_{1,\pi}}\circ F^\circ_{W''_{12},\boldsymbol{\gamma}''_{12},\mathfrak{s}''_{12}}\circ F^\circ_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}.$$

We now use the composition law for path cobordisms (Corollary 10.3) on the 2-handles, and see that this becomes

$$\sum_{\substack{\mathfrak{s}'\in\mathrm{Spin}^c(W'_{22}\circ W'_{21})\\ \mathfrak{s}'|_{W_{2i}}=\mathfrak{s}'_{2i}}} F^\circ_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}}\circ F^\circ_{W''_{31},\boldsymbol{\gamma}'''_{31},\mathfrak{s}''_{31}}\circ F^\circ_{W'_{22}\cup W'_{21},\boldsymbol{\gamma}'_{22}*\boldsymbol{\gamma}''_{21},\mathfrak{s}'}\circ A_{\Gamma''_{2,\pi}*\Gamma''_{1,\pi}}\circ F^\circ_{W''_{12},\boldsymbol{\gamma}''_{12},\mathfrak{s}''_{12}}\circ F^\circ_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}.$$

A Spin^c structure $\mathfrak{s}'$ on $W'_{22}\circ W'_{21}$ uniquely extends over the 1-handles and 3-handles to a Spin structure $\mathfrak{s}$ on $W_2\circ W_1$. It is also easy to see that $\mathfrak{s}|_{W'_{2i}} = \mathfrak{s}'_{2i}$ for $i=1,2$ iff $\mathfrak{s}|_{W_{2i}} = \mathfrak{s}_i|_{W_{2i}}$ for $i=1,2$, which itself occurs iff $\mathfrak{s}|_{W_i} = \mathfrak{s}_i$.

Applying the composition rule for path cobordisms to the 1-handle maps and the 3-handle maps, we thus see that the above expression becomes

$$\sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W)\\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}} F^\circ_{W_{32}\cup W''_{31},\boldsymbol{\gamma}_{32}*\boldsymbol{\gamma}''_{31},\mathfrak{s}|} \circ F^\circ_{W'_{22}\cup W'_{21},\boldsymbol{\gamma}'_{22}*\boldsymbol{\gamma}''_{21},\mathfrak{s}|} \circ A_{\Gamma''_{2,\pi}*\Gamma''_{1,\pi}} \circ F^\circ_{W''_{12}\cup W_{11},\boldsymbol{\gamma}''_{12}*\boldsymbol{\gamma}_{11},\mathfrak{s}|},$$

which is exactly

$$\sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W)\\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}} F^\circ_{W,\Gamma,\mathfrak{s}}$$

by Lemma 12.6. □

13. Cobordisms with $\pi_0(Y_i) \not\twoheadrightarrow \pi_0(W)$

In this section we discuss graph cobordisms which have components with empty ends. The general idea is to remove balls and add paths arbitrarily connecting the new boundary components to the graph whenever a component of W has an empty incoming or outgoing end. The results of this section follow from a technical lemma which we prove first. We then proceed to construct the graph cobordism maps for arbitrary graph cobordisms, and then complete the proof of the graph composition law for arbitrary graph cobordisms. In this section we will finish the proofs of Theorems C and E.

13.1. Main technical lemma.

The following is the main technical lemma of this section:

Lemma 13.1. *Suppose $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ is a graph cobordism with $\pi_0(Y_i)\to\pi_0(W)$ surjective, and suppose that $(W',\Gamma'):(Y_1,\mathbf{w}_1)\to(Y_2\sqcup S^3,\mathbf{w}_2\cup\{w_0\})$ denotes the cobordism obtained by removing a ball from W, and connecting the resulting S^3 arbitrarily with a path to the interior of an edge of Γ. Then*

$$F^\circ_{W',\Gamma',\mathfrak{s}|_{W'}} = F^\circ_{0,\mathfrak{s}|_{Y_2}} \circ F^\circ_{W,\Gamma,\mathfrak{s}}.$$

Similarly if $(W'',\Gamma''):(Y_1\sqcup S^3,\mathbf{w}_1\cup\{w_0\})\to(Y_2,\mathbf{w}_2)$ is the graph cobordism resulting from removing a ball inside of W and connecting the basepoint in S^3 with a path arbitrarily to the interior of an edge of Γ, then

$$F^\circ_{W'',\Gamma'',\mathfrak{s}|_{W''}} = F^\circ_{W,\Gamma,\mathfrak{s}} \circ F^\circ_{4,\mathfrak{s}|_{Y_1\sqcup S^3}},$$

where $F^\circ_{4,\mathfrak{s}|_{Y_1\sqcup S^3}}$.

Proof. We will only prove the statement about the cobordism (W',Γ') since the statement about (W'',Γ'') is proven nearly identically. Using Lemma 12.7 (the composition law) it is sufficient to show that the statement holds for a cobordism of the form $(W_0,\Gamma_0):(Y,\mathbf{w}\cup\{w_0,w_1\})\to(Y,\mathbf{w})$ where $W_0=Y\times I$ and $\Gamma_0=(\mathbf{w}\times I)\cup\widetilde{\lambda}$ where $\widetilde{\lambda}$ is a neatly embedded purturbation of a path λ from w_0 to w_1 in $Y\times\{0\}$. To see that this is sufficient to show the claim for an arbitrary graph cobordism $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ with $\pi_0(Y_i)\to\pi_0(W)$ surjective, note that given any ball $B\subseteq W$ and any choice of path from ∂B to the interior of an edge in Γ, we can isotope that edge, as well as B and the path to the edge, so that they lie in a collar neighborhood of Y_2. Using the composition for such graph cobordisms (Lemma 12.7), we see that showing the claim for (W_0,Γ_0) implies the result for a general graph cobordism (W,Γ).

Let (W'_0,Γ'_0) be the cobordism obtained by removing a ball from W_0 and connecting a basepoint in the new boundary S^3 to a point on $\widetilde{\lambda}$. The cobordisms are depicted in Figure 18.

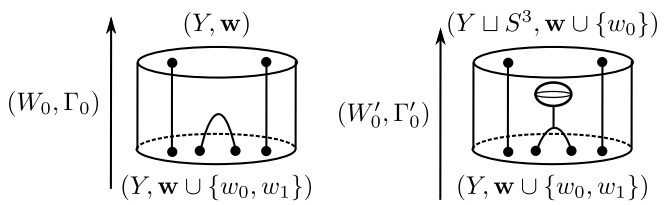

Figure 18. The cobordisms (W_0,Γ_0) and (W'_0,Γ'_0) in Lemma 13.1

We can isotope portions of the graphs Γ_0 and Γ_0' into the incoming ends of their respective cobordisms. Isotope Γ_0 into $Y \times \{0\}$ so that it consists of an arc λ in $Y \times \{0\}$. Isotope Γ_0' in W_0' so that it becomes λ with an additional edge, e added between a point on the interior of λ, to a vertex $(v, 0) \in Y \times \{0\}$, and then a path γ in W_0 from $(v, 0)$ to a point in S^3. We can assume that γ is a subset of $\{v\} \times I$. Let $(W_0', \boldsymbol{\gamma})$ be the path cobordism $(W_0', (\mathbf{w} \times I) \cup \gamma)$. Let $\mathcal{G} = (\lambda, \{w_0, w_1\}, \varnothing, \sigma, \mathbf{w})$ be the flow graph in Y consisting of the path λ. Let $\mathcal{G}'$ be the flow graph $(\lambda \cup e, \{w_0, w_1\}, \{v\}, \sigma, \mathbf{w})$ in $Y \times \{0\}$. By Lemma 12.5, we know that

$$F^\circ_{W_0, \Gamma_0, \mathfrak{s}} = F_{W_0, \mathbf{w} \times I} \circ A_{\mathcal{G}} = A_{\mathcal{G}}$$

and

$$F^\circ_{W_0', \Gamma_0', \mathfrak{s}|_{W_0'}} = F^\circ_{W_0', \boldsymbol{\gamma}, \mathfrak{s}|_{W_0'}} \circ A_{\mathcal{G}'}.$$

Hence the desired equality is equivalent to

$$F^\circ_{0, \mathfrak{s}} \circ A_{\mathcal{G}} = F^\circ_{W_0', \boldsymbol{\gamma}, \mathfrak{s}|_{W_0'}} \circ A_{\mathcal{G}'}.$$

By Lemma 7.11 we know that

$$A_{\mathcal{G}} = S_v^- \circ A_{\mathcal{G}'},$$

and hence it is sufficient to show that

$$F_{W_0', \boldsymbol{\gamma}, \mathfrak{s}|_{W_0'}} = F^\circ_{0, \mathfrak{s}} \circ S_v^-.$$

To see this, put the trivial Morse function $f(y, t) = t$ on $W_0 = Y \times I$. By introducing an index 3-4 critical point pair, and removing a ball B centered at the index 4 critical point, we get a Morse function f for $W_0' = W_0 \setminus B$ which has a single index 3 critical point. A sequence of diagrams encoding the composition is shown in Figure 19.

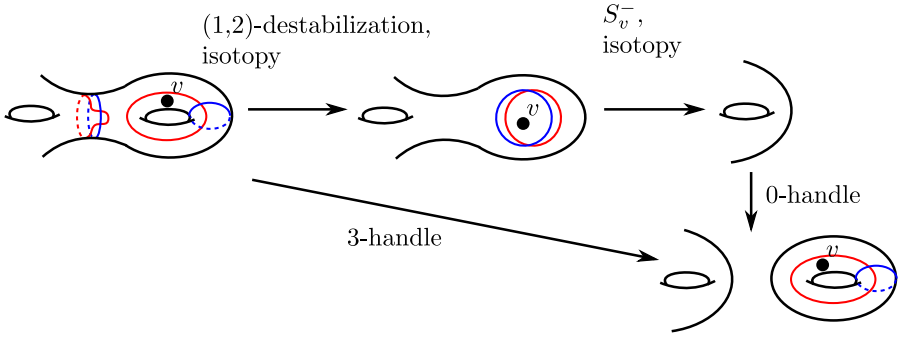

FIGURE 19. A sequence of diagrams which can be used to demonstrate that $F_{W_0', \boldsymbol{\gamma}, \mathfrak{s}|_{W_0'}} = F^\circ_{0, \mathfrak{s}} \circ S_v^-$.

The path consisting of the long arrow down and right is $F_{W_0', \boldsymbol{\gamma}, \mathfrak{s}|_{W_0'}}$. The path consisting of a destabilization, followed by S_v^- followed by the zero handle map is $F^\circ_{0, \mathfrak{s}} \circ S_v^-$. These compositions are easily computed to be equal from the definitions of the maps involved. □

13.2. **Construction of graph cobordism maps for arbitrary cobordisms.** We first define the complexes $CF^\circ(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P})$ by

$$\begin{aligned}
\widehat{CF}(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= \mathbb{Z}_2, \\
CF^-(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= \mathbb{Z}_2[U_{\mathfrak{P}}], \\
CF^\infty(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= \mathbb{Z}_2[U_{\mathfrak{P}}, U_{\mathfrak{P}}^{-1}] \\
CF^+(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= CF^\infty(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) / CF^-(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}),
\end{aligned}$$

all with trivial differentials.

Given an arbitrary cobordism $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$, remove a balls $B_1, \dots, B_k$ from each component of W which doesn't contain a nonempty component of Y_1, and remove additional balls $B_1', \dots, B_n'$ from

each component of W from each component of W which doesn't contain a component of Y_2. Let S_i denote the sphere ∂B_i, and let S_i' denote the sphere $\partial B_i'$. Let

$$(W_0, \Gamma_0, \mathfrak{s}_0) : (Y_1 \sqcup \coprod_i S_i, \mathbf{w}_1 \cup \{w_1, \dots, w_k\}) \to (Y_2 \sqcup \coprod_i S_i', \mathbf{w}_2 \cup \{w_1', \dots, w_n'\})$$

be the graph cobordism where

$$W_0 = W \setminus \left((\bigcup_i B_i) \cup (\bigcup_i B_i') \right).$$

and Γ_0 is the graph obtained by adjoining a path between the basepoint w_i in each S_i and points on Γ on the interior of an edge, and similarly adjoining a path between each of the points $w_i' \in S_i'$ and points on the interior of edges of Γ. The Spin^c structure $\mathfrak{s}_0$ is $\mathfrak{s}|_{W_0}$. We then define the graph cobordism map

$$F^\circ_{W,\Gamma,\mathfrak{s}} \stackrel{\text{def}}{=} \prod_{i=1}^n F^\circ_{4,S_i',\mathfrak{s}|} \circ F_{W_0,\Gamma_0,\mathfrak{s}_0} \circ \prod_{i=1}^k F^\circ_{0,S_i,\mathfrak{s}|},$$

where $\mathfrak{s}|$ denotes the appropriate restriction of $\mathfrak{s}$.

Lemma 13.2. *The map $F^\circ_{W,\Gamma,\mathfrak{s}}$ defined above is independent of the the choice of balls B_i and B_i', as well as the paths from points on the S_i and S_i' to Γ.*

Proof. Suppose we wish to replace one of the balls B_i and the path λ_i from $w_i \in \partial B_i$ to a point on Γ, with another ball $\widetilde{B_i}$ and path $\widetilde{\lambda}_i$. By diffeomorphism invariance of the cobordism map $F_{W_0,\Gamma_0,\mathfrak{s}_0}$, performing an isotopy of B_i or λ_i which fixes Γ leaves the cobordism map unchanged. Hence we can assume that $(B_i \cup \lambda_i)$ and $(\widetilde{B}_i \cup \widetilde{\lambda}_i)$ are disjoint. Let $(W_0, \Gamma_0, \mathfrak{s}_0)$ be the cobordism obtained by removing B_i and adding the path λ_i, and let $(\widetilde{W_0}, \widetilde{\Gamma}_0, \widetilde{\mathfrak{s}}_0)$ be the graph cobordism obtained by removing $\widetilde{B}_i$ instead of B_i ad adding the path $\widetilde{\lambda}_i$. Let $(\overline{W}_0, \overline{\Gamma}_0, \overline{\mathfrak{s}}_0)$ denote the cobordism obtained by removing both B_i and $\widetilde{B_i}$ and adding both paths λ_i and $\widetilde{\lambda}_i$. Applying Lemma 13.1, we see that

$$\begin{aligned} F^\circ_{W_0,\Gamma_0,\mathfrak{s}_0} \circ F^\circ_{0,S_i,\mathfrak{s}|} &= F^\circ_{\overline{W}_0,\overline{\Gamma}_0,\overline{\mathfrak{s}}_0} \circ F^\circ_{0,\widetilde{S}_i,\mathfrak{s}|} \circ F_{0,S_i,\mathfrak{s}|} \\ &= F^\circ_{\overline{W}_0,\overline{\Gamma}_0,\overline{\mathfrak{s}}_0} \circ F_{0,S_i,\mathfrak{s}|} \circ F^\circ_{0,\widetilde{S}_i,\mathfrak{s}|} \\ &= F^\circ_{\widetilde{W}_0,\widetilde{\Gamma}_0,\widetilde{\mathfrak{s}}_0} \circ F^\circ_{0,\widetilde{S}_i,\mathfrak{s}|}, \end{aligned}$$

showing the the replacing B_i and λ_i with $\widetilde{B}_i$ with $\widetilde{\lambda}_i$ doesn't affect the graph cobordism map. The same argument also shows invariance from the B_i' balls and choice of paths. □

We can now prove Theorem C:

Proof of Theorem C. Invariance of the graph cobordism maps for arbitrary graph cobordisms now follows from Lemma 13.2 as well as invariance of the graph cobordism maps for graph cobordisms with $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$ surjective (Theorem 11.1). □

13.3. **Composition law for arbitrary cobordisms.** In this subsection we prove Theorem E, the composition law for arbitrary cobordisms:

Theorem E. *Suppose that $(W_1, \Gamma_1) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ and $(W_2, \Gamma_2) : (Y_2, \mathbf{w}_2) \to (Y_3, \mathbf{w}_3)$ are two graph cobordism maps and $\mathfrak{s}_i \in \text{Spin}^c(W_i)$ are Spin^c structures. Writing (W, Γ) for the composition $(W_2, \Gamma_2) \circ (W_1, \Gamma_1)$, we have the following composition law:*

$$F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s} \in \text{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F^\circ_{W,\Gamma,\mathfrak{s}}.$$

Proof. Pick balls B_{1j} and B_{2j} to remove from W_1 and W_2 respectively so that each component of W_i contains a nonempty component of the incoming boundary. Similarly pick balls B_{1j}' and B_{2j}' in W_1 and W_2 respectively so that each component of W_i contains a nonempty component of the outgoing boundary. Let $S_{ij} = \partial B_{ij}$ and $S_{ij}' = \partial B_{ij}'$. Let W_{i0} be the manifold obtained by removing the balls B_{ij} and B_{ij}'. Pick

paths from points on the new boundary spheres in W_{i0} to Γ_i. Let $\widetilde{\mathbf{w}}_i$ denote the set of basepoints in the spheres S_{ij} and let $\widetilde{\mathbf{w}}'_i$ denote the set of basepoints in the spheres S'_{ij}. Let (W_{i0}, Γ_{i0}) be the resulting graph cobordisms. Define a graph cobordism

$$(\widetilde{W}_{10}, \widetilde{\Gamma}_{10}) : (Y_1 \sqcup_i S_{1i} \sqcup_i S_{2i}, \mathbf{w}_1 \cup \widetilde{\mathbf{w}}_1 \cup \widetilde{\mathbf{w}}_2) \to (Y_2 \sqcup_i S'_{1i} \sqcup_i S_{2i}, \mathbf{w}_2 \cup \widetilde{\mathbf{w}}'_1 \cup \widetilde{\mathbf{w}}_2)$$

by

$$\widetilde{W}_{10} = W_{10} \sqcup_i (S_{2i} \times I)$$

and

$$\widetilde{\Gamma}_{10} = \Gamma_{10} \sqcup (\widetilde{\mathbf{w}}_2 \times I).$$

Similarly define a graph cobordism

$$(\widetilde{W}_{20}, \widetilde{\Gamma}_{20}) : (Y_2 \sqcup_i S_{2i} \sqcup_i S'_{1i}, \mathbf{w}_2 \cup \widetilde{\mathbf{w}}_2 \cup \widetilde{\mathbf{w}}'_1) \to (Y_3 \sqcup_i S'_{2i} \sqcup_i S'_{1i}, \mathbf{w}_3 \cup \widetilde{\mathbf{w}}'_2 \cup \widetilde{\mathbf{w}}'_1)$$

by

$$\widetilde{W}_{20} = W_{20} \sqcup_i (S'_{1i} \times I)$$

and

$$\widetilde{\Gamma}_{20} = \Gamma_{20} \sqcup (\widetilde{\mathbf{w}}'_1 \times I).$$

The cobordisms $(\widetilde{W}_{i0}, \Gamma_{i0})$ are shown in Figure 20.

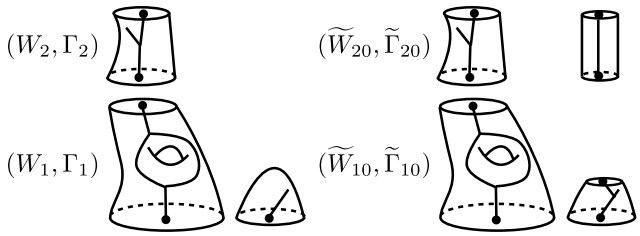

FIGURE 20. The cobordisms (W_i, Γ_i) and $(\widetilde{W}_{i0}, \Gamma_{i0})$, as cobordisms from bottom to top.

We now observe that

$$\begin{aligned} F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1} &= \left(\prod_i F_{4,S'_{2i}}\right) F^\circ_{W_{20},\Gamma_{20},\mathfrak{s}_{20}} \left(\prod_i F^\circ_{0,S_{2i}}\right) \left(\prod_i F^\circ_{4,S'_{1i}}\right) F^\circ_{W_{10},\Gamma_{10},\mathfrak{s}_{10}} \left(\prod_i F^\circ_{0,S_{1i}}\right) \\ &= \left(\prod_i F^\circ_{4,S'_{2i}}\right) \left(\prod_i F^\circ_{4,S'_{1i}}\right) \left(F^\circ_{\widetilde{W}_{20},\widetilde{\Gamma}_{20},\widetilde{\mathfrak{s}}_{20}} \circ F^\circ_{\widetilde{W}_{10},\widetilde{\Gamma}_{10},\widetilde{\mathfrak{s}}_{10}}\right) \left(\prod_i F^\circ_{4,S_{2i}}\right) \left(\prod_i F^\circ_{4,S_{1i}}\right), \end{aligned}$$

where $\mathfrak{s}_{i0} = \mathfrak{s}_i|_{W_{i0}}$ and $\widetilde{\mathfrak{s}}_{i0}$ is the unique Spin^c structure on $\widetilde{W}_{i0}$ which restricts to $\mathfrak{s}_{i0}$ on W_{i0}. The last equality follows since the 0- and 4-handle maps trivially commute with 1-,2- and 3-handle maps, free (de)stabilization maps, and the relative homology maps. Let $(\widetilde{W}_0, \widetilde{\Gamma}_0)$ denote the composition $(\widetilde{W}_{20}, \widetilde{\Gamma}_{20}) \circ (\widetilde{W}_{10}, \widetilde{\Gamma}_{10})$. Using Lemma 12.7, we have that

$$\begin{aligned} &\left(\prod_i F^\circ_{4,S'_{2i}}\right) \left(\prod_i F^\circ_{4,S'_{1i}}\right) \left(F^\circ_{\widetilde{W}_{20},\widetilde{\Gamma}_{20},\widetilde{\mathfrak{s}}_{20}} \circ F^\circ_{\widetilde{W}_{10},\widetilde{\Gamma}_{10},\widetilde{\mathfrak{s}}_{10}}\right) \left(\prod_i F^\circ_{4,S_{2i}}\right) \left(\prod_i F^\circ_{4,S_{1i}}\right) \\ &= \left(\prod_i F^\circ_{4,S'_{2i}}\right) \left(\prod_i F^\circ_{4,S'_{1i}}\right) \left(\sum_{\substack{\mathfrak{s}_0 \in \mathrm{Spin}^c(\widetilde{W}_0) \\ \mathfrak{s}_0|_{\widetilde{W}_{i0}} = \widetilde{\mathfrak{s}}_{i0}}} F^\circ_{\widetilde{W}_0,\widetilde{\Gamma}_0,\mathfrak{s}_0}\right) \left(\prod_i F^\circ_{4,S_{2i}}\right) \left(\prod_i F^\circ_{4,S_{1i}}\right), \end{aligned}$$

which by definition is just

$$\sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(\widetilde{W}) \\ \mathfrak{s}|_{\widetilde{W}_i} = \mathfrak{s}_i}} F^\circ_{W,\Gamma,\mathfrak{s}},$$

since $\widetilde{W}_0$ is obtained by removing balls from W and a Spin^c structure on $\widetilde{W}_0$ uniquely extends over the 4-balls to a Spin^c structure on W. □

14. Moving basepoints, the π_1 action, and the relation between the graph and path cobordism maps

In this section we will prove Theorem D, showing that the graph cobordism maps extend the path cobordism maps when a graph is a collection of paths. We also prove Theorem F and provide an explicit description of the chain homotopy type of the π_1 action. Using this we prove Theorem G, showing that for path cobordisms where each path has a distinct color, the cobordism maps on homology do not depend on the choice of paths.

In addition, given a 3-manifold $(Y, \mathbf{w})$ and a path λ from one of the basepoints $w \in \mathbf{w}$ to a point $w' \notin \mathbf{w}$, we can consider the diffeomorphism of Y corresponding to a finger move along λ. We call the finger move the **basepoint moving map**, and we denote the induced map by $BM(\lambda)$. As an intermediate step to computing the π_1 action, we will see in Lemma 14.7 that the chain homotopy type of the map $BM(\lambda)$ is $S^-_w A_\lambda S^+_{w'}$.

The main technical tool will be a transition map computation similar to [Zem15, Lem. 8.1], though the result we state here for CF^+, CF^∞ and CF^- is substantially harder than the result for $\widehat{CF}$.

14.1. Transition map computation.

The transition map computation that we will need is the one shown in Figure 21.

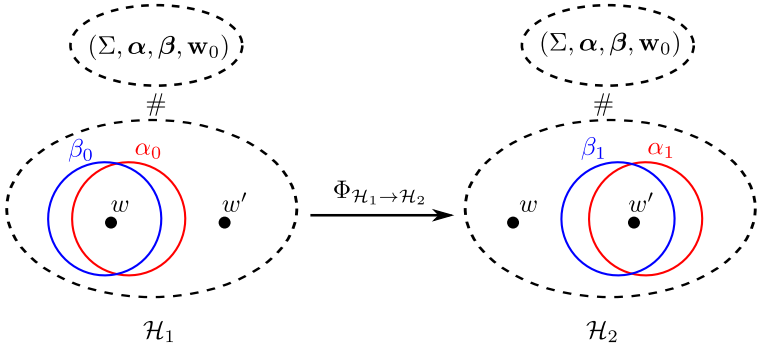

FIGURE 21. We need to compute the transition map pictured above. The case that $\mathbf{w}_0 = \varnothing$ is not excluded.

Theorem 14.1. *For a choice of almost complex structure which can be used to compute $S^\pm_w$ on $\mathcal{H}_1$ and $S^\pm_{w'}$ on $\mathcal{H}_2$, the transition map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ takes the form*

$$\Phi_{\mathcal{H}_1 \to \mathcal{H}_2} = \begin{pmatrix} 1 & 0 \\ \sum_{i,j \geq 0} U_w^i U_{w'}^j \partial^{i+j+1} & 1 \end{pmatrix},$$

where ∂^{i+j+1} is the component of the different $\partial^\infty_{\mathcal{H}_0}$ on the unstabilized diagram $\mathcal{H}_0 = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{z\})$ where z is any point in the domain of the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ containing w and w' and we write

$$\partial^\infty = \sum_{i=0}^{\infty} \partial^i U_z^i.$$

The rest of this subsection is devoted to the proof of Theorem 14.1, though the proof will take several steps. The transition map can be computed as a composition of two triangle maps. We can compute some of the components of each of the corresponding triangle maps explicitly, and the remaining components can be computed by observing that the transition map must be a chain map, and then explicitly computing the differential for several diagrams.

Lemma 14.2. *Let ∂_w and $\partial_{w'}$ be the differentials on the diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\})$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w'\})$ respectively. For sufficiently stretched almost complex structure and for a connected sum point on*

$(S^2, \alpha_i, \beta_i, w, w')$ which is sufficiently close to a point on $\beta_i \setminus \alpha_i$ or $\beta_i \setminus \alpha_i$, the differentials on the diagrams $\mathcal{H}_i$ in Figure 21 are

$$\partial_{\mathcal{H}_1} = \begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix}$$

and

$$\partial_{\mathcal{H}_2} = \begin{pmatrix} \partial_{w} & U_w + U_{w'} \\ 0 & \partial_{w} \end{pmatrix}.$$

Proof. This is [OS08, Proposition 6.5] in the case that the connected sum point approaches a point on $\beta_0 \setminus \alpha_0$. The case that the connected sum point approaches a point on $\alpha_0 \setminus \beta_0$ is identical. □

We also need the differential of the diagram shown in Figure 22.

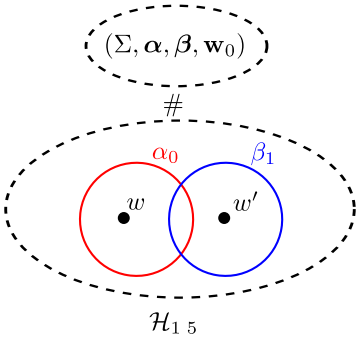

FIGURE 22. The intermediate diagram $\mathcal{H}_{1.5}$.

Lemma 14.3. *The diagram $\mathcal{H}_{1.5}$, shown in Figure 22, is strongly $\mathfrak{s}$-admissible if $\mathcal{H}_0$ is strongly $\mathfrak{s}$-admissible, and the differential is*

$$\partial_{\mathcal{H}_{1.5}} = \begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix}$$

for sufficiently stretched almost complex structure and connected sum point on $(S^2, \alpha_1, \beta_0, w, w')$ taken sufficiently close to a point on $\beta_1 \setminus \alpha_0$.

Proof. This will follow from a careful analysis of the proof of [OS08, Prop. 6.5], as well as a degeneration argument. We should note that this does not follow immediately from the analysis in [OS08].

Some multiplicities for a disk on $(S^2, \alpha_1, \beta_0, w, w')$ are shown in Figure 23.

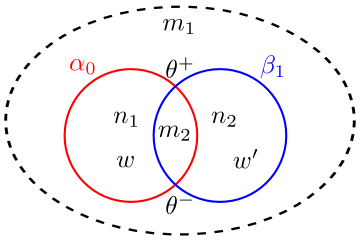

FIGURE 23. Multiplicities on the diagram $(S^2, \alpha_1, \beta_0, w, w')$ used to analyze disks in $\mathcal{H}_{1.5}$.

Write

$$\partial_{\mathcal{H}_{1.5}} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

We now compute A and D. Suppose

$$\phi = \phi_1 \# \phi_2 \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^+) \qquad \text{or} \qquad \phi \in \pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^-).$$

By the same Maslov index computation as in [Zem15, Lemma 8.3], we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 = \mu(\phi_1) + 2m_2.$$

From positivity, we conclude that for sufficiently stretched neck that $\mu(\phi_1) = 1$ and $m_2 = 0$. By the vertex relations we also have that $n_1 + n_2 = m_1$.

In Lemma 6.5 we performed a similar count of disks, following the analysis in [OS08, Prop. 6.5]. We showed that given a disk on the unstabilized diagram $\mathcal{H}_0 = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\})$, there is one (modulo two) Maslov index $2m_1$ disk in $\pi_2(\theta^{\pm}, \theta^{\pm})$ which satisfies the matching condition. We thus have an identification of the analogous components of $\partial_{\mathcal{H}_1}$ as $\partial_{\mathcal{H}_0}$. On the other hand, in the diagram $\mathcal{H}_{1.5}$, the disks can cross over either w and w', and we only have the relation $n_w(\phi) + n_{w'}(\phi) = m_1$. Hence we can write

$$A = \sum_{i,j \geq 0} m_{ij} U_w^i U_{w'}^j$$

where m_{ij} are maps such that

$$\sum_{i+j=k} m_{ij} = \partial^k_{\mathcal{H}_0}.$$

Note that ϕ_2 is homologous to a sum of α_0 and β_1 boundary degenerations with $m_2 = 0$. We claim that if we degenerate the almost complex structure by letting the connected sum point on S^2 approach a point in $\beta_1 \setminus \alpha_0$, then ϕ_2 can consist only of α_0 boundary degenerations. If we establish this, then we will have that $n_1 = 0$ and $n_2 = m_1$, which would establish $A = D = \partial_{w'}$.

Let $p_2^n \in S^2$ denote a sequence of connected sum points approaching a point $p_2^\infty \in \beta_1 \setminus \alpha_0$. Suppose that $\phi = \phi_1 \# \phi_2$ is a homology class with holomorphic representatives u_T^n for all arbitrarily large n and all large neck lengths T. For each n, by sending $T \to \infty$, we get that a subsequence of the u_T^n converges to a broken holomorphic disk U^n on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ of total homology ϕ_1 and a broken holomorphic disk V^n on (S^2, α_0, β_1) of total homology ϕ_2. Since $\mu(\phi_1) = 1$ we conclude that U^n can consist of only a single holomorphic strip u^n of Maslov index 1. There must be some component v_1^n in V which satisfies the matching condition with respect to u^n, i.e.

$$\rho^{p_1}(u^n) = \rho^{p_2^n}(v_1^n).$$

We now claim that any other components of V^n must have Maslov index zero. Let ϕ_2' denote the homology class of v_1^n. By the matching condition, we know that $\phi_1 \# \phi_2'$ admits a holomorphic representative, $u^n \# v_1^n$. By transversality we conclude that $\mu(\phi_1 \# \phi_2') \geq 1$. Let the remaining components of V^n be $v_2^n, \dots, v_k^n$ and say that v_2^n represents the homology class ϕ_2^i for $i = 2, \dots, k$. Note that due to the matching condition we have that $m_1(\phi_2^i) = 0$. Since each ϕ_2^i admits a holomorphic representative and is either a holomorphic strip, or a boundary degeneration (or in principle a closed surface, though these are excluded by $m_1(\phi_2^i) = 0$), we know that $\mu(\phi_2^i) > 0$ iff $\phi_2^i \neq 0$. But we have

$$1 = \mu(\phi_1 \# \phi_2) = \mu(\phi_1 \# \phi_2' + 0 \# \phi_2^2 + \cdots + 0 \# \phi_2^k) = \mu(\phi_1 \# \phi_2') + \mu(0 \# \phi_2^2) + \cdots + \mu(0 \# \phi_2^k)$$

$$= \mu(\phi_1 \# \phi_2') + \mu(\phi_2^2) + \cdots + \mu(\phi_2^k)$$

By our previous observations the only way for this equality to be true is for $\mu(\phi_1 \# \phi_2') = 1$ and $\mu(\phi_2^i) = 0$ for all i. Hence v_1^n is the only nonconstant component of V^n. Let v^n denote v_1^n.

We now let $p^n \to p^\infty$. Note that u^n is a holormorphic representative for ϕ_1 and $\mu(\phi_1) = 1$. Modulo the $\mathbb{R}$-action, there are only finitely many such disks. Since there are only finitely many representatives of ϕ_1 modulo translation, we can simply assume that $\rho^{p_1}(u^n) = \Delta$ for some fixed divisor $\Delta \subseteq \mathrm{Sym}^{m_1}(\mathbb{D})$. The matching condition then reads that

$$\rho^{p_2^n}(v^n) = \Delta.$$

The sequence v^n has a Gromov convergent subsequence to some broken holomorphic disk V^∞ on $(S^2, \alpha_0, \beta_1, w, w')$. Since $\rho^{p_2^n}(v^n) = \Delta$, there must be some component v_1^∞ of V^∞ with $\Delta \subseteq \rho^{p_2^\infty}(v_1^\infty)$. Also note that $|\Delta| = |m_1(\phi)|$. Hence there are $|\Delta| = |m_1(\phi)|$ points in the domain of v_1^∞ which map to p_2^∞ and since Δ doesn't include any points on $\{0,1\} \times \mathbb{R} \subseteq \mathbb{D}$, we know that these points are on the interior of the domain of v_1^∞. By the open mapping theorem applied to $\pi_\Sigma \circ v_1^\infty$, we conclude that

$$m_1(v_1^\infty) \geq |\Delta| = m_1(\phi_2)$$

and

$$n_2(v_1^\infty) \geq |\Delta| = m_1(\phi_2).$$

By positivity, the remaining components can only increase the multiplicity in each region, so we thus conclude that

$$n_2(\phi_2) \geq m_1(\phi_2).$$

Strictly speaking, the open mapping theorem doesn't apply since $\pi_\Sigma \circ v_1^\infty$ is not a holomorphic map, but one can use [MW95] to show that $\pi_\Sigma \circ v_1^\infty$ is an open map on the interior of its domain (cf. [Lip06, Lem. 3.1.]). On the other hand, through the Maslov index computation we already knew that $n_2(\phi_2) \leq m_1(\phi_2)$ so we conclude that $n_2(\phi_2) = m_1(\phi_2)$. Since $\phi_2 \in \pi_2(\theta^+, \theta^+)$ or $\phi_2 \in \pi_2(\theta^-, \theta^-)$, we recall that our previous computation showed that $m_2(\phi_2) = 0$ so that

$$n_1(\phi_2) + n_2(\phi_2) = m_1(\phi_2),$$

and hence we conclude that $n_1(\phi_1) = 0$ and $n_2(\phi_1) = m_1(\phi_2)$, concluding the computation.

The computation of the remaining components, B and C, follows nearly identically to [OS08, Proposition 6.5], the only caveat being that we basepoints are in slightly different locations than in [OS08, Proposition 6.5], but the disks can be counted the same and the result is the same. The other difference from [OS08, Proposition 6.5] is that we are letting the connected sum point approach a point in $\alpha_1 \setminus \beta_0$ as opposed to $\beta_0 \setminus \alpha_1$, but their argument works essentially without change. □

Having computed the differentials for the diagrams $\mathcal{H}_1, \mathcal{H}_{1.5}$ and $\mathcal{H}_2$, we now proceed to compute the transition map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$. Observe that the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2} = \Phi_{\mathcal{H}_{1.5} \to \mathcal{H}_2} \circ \Phi_{\mathcal{H}_1 \to \mathcal{H}_{1.5}}$ and that each of $\Phi_{\mathcal{H}_{1.5} \to \mathcal{H}_2}$ and $\Phi_{\mathcal{H}_1 \to \mathcal{H}_{1.5}}$ can be computed by a triangle map.

Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{z\})$ denote the diagram with neither the w nor w' basepoint and no α_0 or β_0 curve and instead a single basepoint, z, in the region containing w and w'. Let $\boldsymbol{\alpha'}$ and $\boldsymbol{\beta'}$ be small Hamiltonian isotopies of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ respectively such that

$$|\alpha_i \cap \alpha_j'| = |\beta_i \cap \beta_j'| = 2\delta_{ij}$$

for all i and j.

We now consider the triples

$$\mathcal{T}_1 = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\beta'} \cup \{\beta_1\}, \mathbf{w}_0 \cup \{w, w'\})$$

and

$$\mathcal{T}_2 = (\Sigma, \boldsymbol{\alpha'} \cup \{\alpha_1\}, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta'} \cup \{\beta_1\}, \mathbf{w}_0 \cup \{w, w'\})$$

depicted in Figure 24 and Figure 24 respectively.

The hypotheses of the following lemma may seem somewhat strange, but we will later see how to satisfy them.

Lemma 14.4. *Suppose there is a sequence of connected sum points $p_2^n \in S^2$ approaching a point $p_2^\infty \in \beta_0 \setminus (\alpha_0 \cup \beta_1)$ such that the differentials of $\partial_{\mathcal{H}_1}$ and $\partial_{\mathcal{H}_{1.5}}$ are as in Lemmas 14.2 and 14.2. Then for n sufficiently large, the triangle map $F_{\mathcal{T}_1}$ takes the form*

$$F_{\mathcal{T}_1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Proof. Write the triangle map as

$$F_{\mathcal{T}_1} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

If we can show that $A = 1$ and $B = 0$, then using the fact that $F_{\mathcal{T}_1}$ is a chain map and applying the hypotheses that the differentials $\partial_{\mathcal{H}_1}$ and $\partial_{\mathcal{H}_{1.5}}$ are as in Lemmas 14.2 and 14.3, we see that

$$\begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ C & D \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ C & D \end{pmatrix} \begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix}.$$

This immediately implies that $D = 1$ and $C(U_w + U_{w'}) = 0$, implying $C = 0$. Hence we just need to show that $A = 1$ and $B = 0$. To do this, we will explicitly count triangles.

We begin with the computation for A. Suppose that $\psi = \psi_1 \# \psi_2$ where $\psi_2 \in \pi_2(x^+, \theta^+, y^+)$. Let m_1 and m_2 be the multiplicities shown in Figure 24.

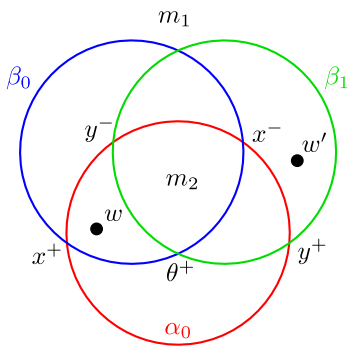

FIGURE 24. Multiplicities for a triangle on $(S^2, \alpha_0, \beta_0, \beta_1)$. These are counted by the transition map in Lemma 14.4.

By a similar argument as in Lemma 6.18 we have that $\mu(\psi_2) = 2m_1 + 2m_2$. Hence

$$\mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1) + 2m_2.$$

For sufficiently stretched neck, by transversality we know that $\mu(\psi_1) \geq 0$. From positivity, we also have that $m_2 \geq 0$. We conclude that $\mu(\psi_1) = 2m_2 = 0$.

Given a sequence of holomorphic triangles u_{T_i} representing ψ with the complex structure $J_s(T_i)$ for a sequence of T_i approaching ∞ as $i \to \infty$ which achieve transversality, by a Gromov compactness argument, we thus get that a subsequence of the u_{T_i} converges to a pair U_1, U_2 where U_i is a broken holomorphic triangle representing ψ_i. Since $\mu(\psi_1) = 0$, we conclude that U_1 can have only one component, namely a single holomorphic triangle u_1. Now U_2 must have a component u_2 with the matching condition $\rho^{p_1}(u_1) = \rho^{p_2}(u_2)$ where p_i are the points where the connected sum is taken. By the same argument as in Theorem 6.17 we conclude that the remaining components of U_2 must have Maslov index zero. Since they consist of holomorphic flowlines, boundary degenerations, and closed surfaces with total Maslov index zero, we conclude that u_2 is the only component of U_2. Hence u_{T_i} limits to a preglued holomorphic triangle.

On the other hand, by Lemma 6.19, given a holomorphic triangle u_1 representing ψ_1, there is one (modulo two) holomorphic triangle u_2 representing a triangle in $\pi_2(x^+, \theta^+, y^+)$ satisfying the matching condition with u_1. Hence we conclude that $A = 1$.

We now need to show that $B = 0$ for certain complex structures. This computation is in two steps. First we perform a Maslov index computation, and then we degenerate by allowing the connected sum point $p_2 \in S^2$ to approach a point in $\beta_0 \setminus (\alpha_0 \cup \beta_1)$.

Write $\psi = \psi_1 \# \psi_2$ for $\psi_1 \in \pi_2(\mathbf{x}, \Theta^+, \mathbf{y})$ and $\psi_2 \in \pi_2(x^-, \theta^+, y^+)$. An argument similar to Lemma 6.18 shows that $\mu(\psi_2) = 2m_1 + 2m_2 - 1$. Hence we have that

$$\mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1) + 2m_2 - 1.$$

By stretching the neck sufficiently, we can assume that $\mu(\psi_1) \geq 0$. Also $m_2 \geq 0$ by positivity, and hence we conclude that $\mu(\psi_1) = 1$ and $m_2 = 0$.

Multiplicities for a $\psi_2 \in \pi_2(x^-, \theta^+, y^+)$ are determined by four integers, m_1, a, b and c, as shown in Figure 25.

We now claim that for p_2 sufficiently close to $\beta_0 \setminus (\alpha_0 \cup \beta_1)$, we must have $a \geq 0$. Assuming that $a \geq 0$, let's first see how this implies that $B = 0$. Notice that we have that

$$m_2 = m_1 + a + b + c = 0.$$

On the other hand, by positivity we also have

$$m_1 + b + c - 1 \geq 0.$$

Since $a \geq 0$, adding a to the above expression and rearranging, we thus have that

$$m_2 = m_1 + b + c + a \geq 1,$$

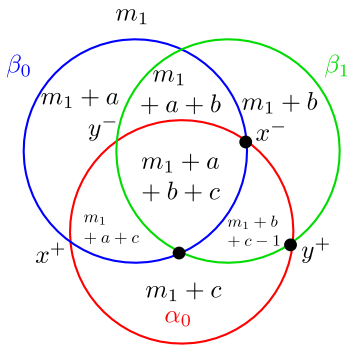

FIGURE 25. Multiplicities for a homology triangle $\psi_2 \in \pi_2(x^-, \theta^+, y^+)$. These are used to compute B in Lemma 14.4.

a contradiction.

We now show that if we let p_2 approach $\beta_0 \setminus (\alpha_0 \cup \beta_1)$, then $a \geq 0$. This argument is similar to the argument we made in Lemma 14.3.

Let p_2^n be a sequence of points in the region labelled by m_1 in Figure 25 which approach a point $p_2^\infty \in \beta_0 \setminus (\alpha_0 \cup \beta_1)$. Suppose there are holomorphic representatives $u_{T_i}^n$ for all large n and arbitrarily large neck lengths T_i. By letting $T_i \to \infty$, the holomorphic triangles $u_{T_i}^n$ converge to a pair (U^n, V^n) where U^n is a broken holomorphic triangle representing ψ_1 and V^n is a broken holomorphic triangle representing ψ_2.

Since $\mu(\psi_1) = 1$, we know that U^n must consist of either a single Maslov index 1 holomorphic triangle, or a Maslov index 0 triangle and a Maslov index 1 holomorphic strip. We will now show that the only possibility is that U^n consists of a single Maslov index 1 triangle.

Suppose to the contrary that U^n consists of a Maslov index 1 disk u_1^n and a Maslov index 0 triangle u_2^n. Then u_1^n and u_2^n must both have matching components v_1^n and v_2^n, i.e. a holomorphic strip v_1^n and a holomorphic triangle v_2^n in V^n with

$$\rho^{p_1}(u_i^n) = \rho^{p_2}(v_i^n).$$

Using the standard gluing techniques, we can glue to form holomorphic curves $u_1^n \# v_1^n$ and $u_2^n \# v_2^n$. By transversality we must have that

$$\mu(u_1^n \# v_1^n) \geq 1$$

and

$$\mu(u_2^n \# v_2^n) \geq 0.$$

If τ is any of the remaining components of V^n, then since U^n has no nonconstant components other than u_1^n and u_2^n, we know that $m_1(\tau) = 0$ and hence τ naturally determines a holomorphic curve in the homology class $0 \# \mathcal{D}(\tau)$ on $\mathcal{T}_1$. In particular, the Maslov index of any of the remaining curves τ must be nonnegative, and hence can only increase the total Maslov index. Since we already have that

$$\mu(u_1^n \# v_1^n) + \mu(u_2^n \# v_2^n) \geq 1,$$

we thus know that $\mu(\psi_1 \# \psi_2) \geq 1$, a contradiction.

We can thus conclude that U^n consists of a single holomorphic triangle u^n of Maslov index 1 and that V^n is a broken holomorphic triangle containing a genuine holomorphic triangle v_1^n which satisfies

$$\rho^{p_1}(u^n) = \rho^{p_2^n}(v_1^n).$$

We now claim that there are no more components of V^n. This proceeds in a similar spirit to the previous argument. Let ψ_2' denote the homology class of the triangle v_1^n. Then $\psi_2 - \psi_2'$ represents the remaining components of V^n. Since the homology triangle $\psi_1 \# \psi_2'$ admits a holomorphic representative $u^n \# v_1^n$, we know that $\mu(\psi_1 \# \psi_2') \geq 0$. Note that because of the matching condition, we have that $m_1(\psi_2 - \psi_2') = 0$. Let $v_2^n, \ldots, v_k^n$ denote the remaining curves of V^n. Let $\phi_2, \ldots, \phi_k$ denote the homology classes of these curves. The homology classes $\phi_2, \ldots, \phi_k$ could consist of disks and boundary degenerations, and closed surfaces. Since these curves have holomorphic representatives and are of the above types, we know that $\mu(\phi_i) > 0$ iff

ϕ_i is nonzero. Note that $n_{p_2^n}(\phi_i) = 0$ for $i = 2, \dots, k$. Note that $\psi_1 \# \psi_2$ is obtained by splicing in the curves $0 \# \phi_i$. An easy computation shows that $\mu(0\#\phi_i) = \mu(\phi_i)$. Hence we have that

$$\mu(\psi_1 \# \psi_2) = \mu(\psi_1 \# \psi_2') + \mu(\phi_2) + \cdots \mu(\phi_k).$$

Since $\mu(\psi_1 \# \psi_2') \geq 0$, and $\mu(\psi_1 \# \psi_2) = 0$ by assumption, we conclude that

$$\mu(\psi_1 \# \psi_2') = \mu(\phi_2) = \cdots = \mu(\phi_k) = 0.$$

In particular we know that the curves $\phi_2, \dots, \phi_k$ can only be constant disks.

Hence V^n consists of a single holomorphic triangle satisfying the matching condition with u^n.

We now let the connected sum point p_2^n approach a point $p_2^\infty \in \beta_0 \setminus (\alpha_0 \cup \beta_1)$. For each n, the previous argument shows that there is a pair (u^n, v^n) of holomorphic triangles with $\rho^{p_1}(u^n) = \rho^{p_2^n}(v^n)$. We can extract Gromov limits of u^n and v^n to broken holomorphic triangles U and V on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta'})$ and $(S^2, \alpha_0, \beta_0, \beta_1)$ respectively.

As before, U either consists of a single Maslov index 1 triangle, or a Maslov index 0 triangle and a Maslov index 1 disk. In either case, the argument proceeds similarly. In the case that U consists of single Maslov index 1 triangle u (resp. a Maslov index 1 disk u' and a Maslov index 0 triangle u'') there must be a matching component $v \in V$ (resp. two matching components v' and v'' in V) with

$$\rho^{p_1}(u) \subseteq \rho^{p_2^\infty}(v)$$

(resp. $\rho^{p_1}(u') \subseteq \rho^{p_2^\infty}(v')$ and $\rho^{p_1}(u'') \subseteq \rho^{p_2^\infty}(v'')$). Arguing using the open mapping theorem (or more precisely using [MW95] or [Lip06, Lem. 3.1.], as we did in Lemma 14.3), if U consists of a single triangle u, then we get that

$$m_1(\phi_1) + a \geq |\rho^{p_1}(u)| = m_1(\phi_1)$$

and hence $a \geq 0$. In the case where U consists of a disk and a triangle, we similarly have $a \geq 0$. □

We now use the same strategy to analyze the second triangle map $F_{\mathcal{T}_2}$ which computes the transition map $\Phi_{\mathcal{H}_{1.5} \to \mathcal{H}_2}$. The Heegaard triple $\mathcal{T}_2$ is depicted in Figure 26.

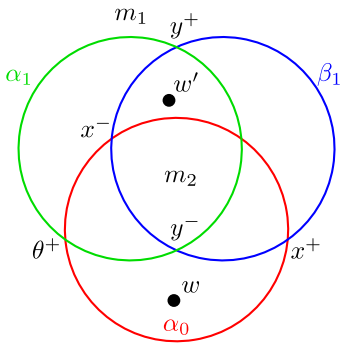

FIGURE 26. Part of the Heegaard triple $\mathcal{T}_2$. This is used to compute $\Phi_{\mathcal{H}_{1.5} \to \mathcal{H}_2}$ in Lemma 14.5.

Lemma 14.5. *Suppose there is a sequence of points p_2^n approaching a point $p_2^\infty \in \alpha_1 \setminus (\alpha_0 \cup \beta_1)$ such that the differentials of $\partial_{\mathcal{H}_2}$ and $\partial_{\mathcal{H}_{1.5}}$ are as in Lemmas 14.2 and 14.2. For n sufficiently large, the transition map takes the form*

$$F_{\mathcal{T}_2} = \begin{pmatrix} 1 & 0 \\ \sum_{i,j \geq 0} U_w^i U_{w'}^j \partial^{i+j+1} & 1 \end{pmatrix}.$$

Proof. The computation is very similar to Lemma 14.4. Write the triangle map

$$F_{\mathcal{T}_1} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

We will show that $B = 0$ and $D = 1$. To compute the other components, observe that by assumption the differentials are of the forms stated in Lemmas 14.2 and 14.3. Since $F_{\mathcal{T}_2}$ is a chain map, we see that

$$\begin{pmatrix} A & 0 \\ C & 1 \end{pmatrix} \begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix} = \begin{pmatrix} \partial_w & U_w + U_{w'} \\ 0 & \partial_w \end{pmatrix} \begin{pmatrix} A & 0 \\ C & 1 \end{pmatrix}.$$

By multiplying out these matrices, we see that $A(U_w + U_{w'}) = U_w + U_{w'}$, implying that $A = 1$. Similarly we have that $C(U_w + U_{w'}) = \partial_w + \partial_{w'}$, implying that $C = \sum_{i,j \geq 0} U_w^i U_{w'}^j \partial^{i+j+1}$ by an easy algebraic computation

We can compute D directly, using the same technique we used to compute the A term in Lemma 14.4. To compute B we will have to employ the same degeneration technique as before.

We now compute D. Suppose that $\psi = \psi_1 \# \psi_2$ is a homology triangle where ψ_1 is a homology triangle in $\pi_2(\mathbf{x}, \Theta^+, \mathbf{y})$ and ψ_2 is a homology triangle in $\pi_2(x^-, \theta^+, y^-)$. Suppose that $\mu(\psi) = 0$. Letting m_1 and m_2 be the multiplicities as in Figure 26, an argument similar to [Zem15, Lemma 8.3] shows that $\mu(\psi^2) = 2m_1 + 2m_2$, and hence

$$\mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1) + 2m_2.$$

By stretching the neck sufficiently, we can conclude that $\mu(\psi_1) = 2m_2 = 0$. Given a holomorphic representative u_{T_i} for ψ for the almost complex structure $J_s(T_i)$ for all i, we can extract a limit as a pair of broken holomorphic triangles U on $(\Sigma, \boldsymbol{\alpha'}, \boldsymbol{\alpha}, \boldsymbol{\beta'})$ and V on $(S^2, \alpha_1, \alpha_0, \beta_1)$. Using the same reasoning as in Lemma 14.4 shows that U and V both consist of exactly one holomorphic triangle each and that the pair are preglued. Conversely given such a preglued pair, we can glue to construct a holomorphic triangle for sufficiently stretched neck. Using the same logic as in Lemma 14.4 we conclude that $D = 1$.

We now wish to show that $B = 0$. We first perform a Maslov index calculation. If $\psi_2 \in \pi_2(x^-, \theta^+, y^+)$ is a homology triangle on $(S^2, \alpha_1, \alpha_0, \beta_1, w, w')$ as in Figure 26, then the Maslov index is

$$\mu(\psi_2) = 2m_1 + 2m_2 - 1.$$

This is proven by verifying the formula for the Maslov index 1 triangle in $\pi_2(x^-, \theta^+, y^+)$ which is formed by splicing a Maslov index one disk from y^- to y^+ with a the "trivial" Maslov index zero triangle in $\pi_2(x^-, \theta^+, y^-)$. Any other triangle differs by splicing in boundary degenerations, and the formula respects such an operation.

Hence we have that if $\psi = \psi_1 \# \psi_2$ then

$$\mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1) + 2m_2 - 1,$$

from which we conclude that if ψ has holomorphic representatives for arbitrarily large neck lengths, then $\mu(\psi_1) = 1$ and $m_2 = 0$.

An arbitrary homology triangle ψ_2 is shown in Figure 27.

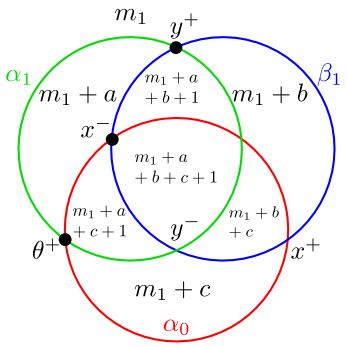

FIGURE 27. Multiplicities for a homology triangle $\psi_2 \in \pi_2(x^-, \theta+, y^+)$ used to compute B in Lemma 14.5.

We claim that by sending the degeneration point p_2 to a point on $\alpha_1 \setminus (\beta_1 \cup \alpha_0)$, we can ensure $a \geq 0$. The argument is identical to the argument in Lemma 14.4, so we omit the argument.

Assuming that $a \geq 0$, we observe that $m_2 = m_1 + a + b + c + 1 = 0$ by the Maslov index calculation. On the other hand we have that $m_1 + b + c \geq 0$ by positivity. Hence $m_1 + a + c + b \geq 0$ and hence

$$m_1 + a + b + c + 1 \geq 1,$$

a contradiction since $m_1 + a + b + c + 1 = 0$. So no triangles are counted and $B = 0$. Hence under the assumptions of the Lemma, the transition map takes the desired form. □

We are now in position to prove Theorem 14.1.

Proof of Theorem 14.1. Notice that if we can satisfy the hypotheses of Lemmas 14.4 and 14.5, then the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ can simply be computed by

$$\Phi_{\mathcal{H}_{1.5} \to \mathcal{H}_2} \circ \Phi_{\mathcal{H}_1 \to \mathcal{H}_{1.5}} = \begin{pmatrix} 1 & 0 \\ \sum_{i,j \geq 0} U_w^i U_{w'}^j \partial^{i+j+1} & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \sum_{i,j \geq 0} U_w^i U_{w'}^j \partial^{i+j+1} & 1 \end{pmatrix}.$$

The issue is simply that we need to be careful about placing the point p_2 in relation to the curves $\alpha_0, \beta_0, \alpha_1$, and β_1.

It is easiest to do this in stages. First place the α_0, β_0 and β_1 curves arbitrarily so that they have the correct configuration for $\mathcal{T}_1$. As in Figure 28, there is an open set R_1 where we can put p_2 so that the differential $\partial_{\mathcal{H}_1}$ is correct and so the degeneration argument in Lemma 14.4 works. Similarly there is an open set R_2 where we can put p_2 so that the differential $\partial_{\mathcal{H}_{1.5}}$ is correct. If the point p_2 is in the set $R_1 \cap R_2$, then $F_{\mathcal{T}_1}$ will thus have the form stated in Lemma 14.4.

We observe that $R_1 \cap R_2$ is nonempty, as is shown in Figure 28. From here, we place the α_1 curve so that there is a point in $\alpha_1 \setminus (\alpha_0 \cup \beta_0 \cup \beta_1)$ which lies in $R_1 \cap R_2$. We simply pick a sequence of points in $R_1 \cap R_2$ which approach a point on $\alpha_1 \setminus (\alpha_0 \cup \beta_0 \cup \beta_1)$ and which is also in $R_1 \cap R_2$. This ensures that $\partial_{\mathcal{H}_2}$ is of the form of Lemma 14.2, and that the degeneration argument in Lemma 14.5 goes through.

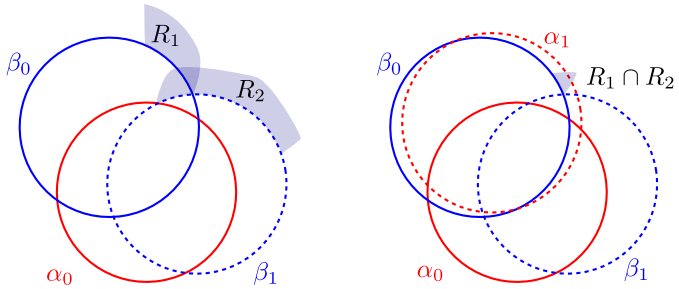

FIGURE 28. Positioning the curves $\alpha_0, \beta_0, \alpha_1, \beta_1$ and the connected sum point. We do this in stages to ensure the requirements of the various lemmas are satisfied.

□

14.2. **An easier model computation.** In this subsection we perform a model computation that we needed to prove various results about the graph cobordism maps. The following implies the result of Lemma 7.7, and will also be used later to compute the $\pi_1(Y, w)$-action.

Lemma 14.6. *Let $\mathcal{H}_1$ be the diagram in the previous section and let λ be an arc between w and w', as in Figure 29. In the almost complex structure used in Theorem 14.1, we have the following chain homotopy equivlence:*

$$A_\lambda \simeq \begin{pmatrix} 0 & U_{w'} \\ 1 & 0 \end{pmatrix}.$$

Proof. This follow relatively immediately from the same analysis one uses to prove Lemma 14.3. We sketch the argument. We first write

$$A_\lambda = \begin{pmatrix} a_\lambda & b_\lambda \\ c_\lambda & d_\lambda \end{pmatrix}.$$

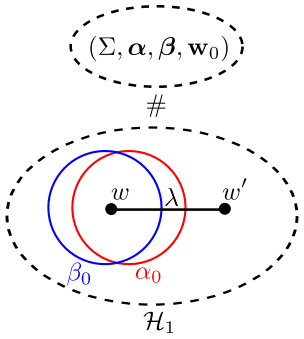

FIGURE 29. The curve λ in Lemma 14.6.

As in Lemma 14.3 by sending the connected sum point towards a point on $\beta_0 \setminus \alpha_0$, we ensure that any disks counted by a_λ or d_λ are disks on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0)$ glued to a disk on $(S^2, \alpha_0, \beta_0, w, w')$ which is homologous to a β_0 degeneration. This immediately forces $a_\lambda = d_\lambda = 0$.

Computing c_λ is straightforward: there are only two disks counted by those maps, namely the two small strips between α_0 and β_0 which don't go over w. Counting these with multiplicities induced by $a(\lambda, \phi)$ we get that $c_\lambda = 1$.

We finally compute b_λ. Following the proof of [OS08, Prop. 6.5] we observe that for sufficiently stretched neck length, and for connected sum point sufficiently close to $\beta_0 \setminus \alpha_0$, we know that there are two domains which contribute disks to b_λ and b_γ. To describe these, let ϕ_w and $\phi_{w'}$ be the two Maslov index 1 disks on $(S^2, \alpha_0, \beta_0, w, w')$ which go over the w basepoint (resp. the w' basepoint) exactly once. Homology disks with domain $0\#\mathcal{D}(\phi_w)$ contribute 0 to b_λ.

The other homology disks which contribute to b_λ have domain equal to $\mathcal{D}\#\mathcal{D}(\phi_{w'})$, where $\mathcal{D}$ is a boundary degeneration if $|\mathbf{w}_0| > 0$ and $\mathcal{D} = [\Sigma]$ if $|\mathbf{w}_0| = 0$. Such a domains represent homology classes of disks in $\pi_2(\mathbf{x} \times \theta^-, \mathbf{x} \times \theta^+)$ for any $\mathbf{x}$. Following the proof of [OS08, Prop. 6.5] for a given $\mathbf{x}$, by sending the connected sum point to $\beta_0 \setminus \alpha_0$ and stretching the neck, there is one (mod 2) holomorphic disk representing a class in $\pi_2(\mathbf{x} \times \theta^-, \mathbf{x} \times \theta^+)$ with domain $\mathcal{D}\#\mathcal{D}(\phi_{w'})$. Such disks contribute $U_{w'}$ to b_λ. □

14.3. **Basepoint moving map, the $\pi_1(Y, p)$-action, and the relation between path and graph cobordism maps.** In this section we compute the $\pi_1(Y, p)$ action on $CF^\circ(Y, p, \mathfrak{s})$. The reader is encouraged to look at [Sar11b], where similar maps on various flavors on knot Floer homology are constructed, but with the homology action maps replaced by the map on knot Floer homology which counts disks which go over the z basepoint exactly once.

Suppose that $\mathbf{w}_0$ is a (possibly empty) collection of basepoints and w and w' are distinct basepoints which aren't in $\mathbf{w}_0$. Given an embedded path λ from w to w', we can consider the diffeomorphism $\phi : (Y, \mathbf{w}_0 \cup \{w\}) \to (Y, \mathbf{w}_0 \cup \{w'\})$ resulting from a finger move along the curve λ. We let $BM(\lambda)$ denote the map induced on the $\mathbb{Z}_2[U_{\mathbf{w}}]$-equivariant chain homotopy type

$$BM(\lambda) = \phi_* : CF^\circ(Y, \mathbf{w}_0 \cup \{w\}, \mathfrak{s}) \to CF^\circ(Y, \mathbf{w}_0 \cup \{w'\}, \mathfrak{s}).$$

Lemma 14.7. *The $\mathbb{Z}_2[U_{\mathfrak{P}}]$-chain homotopy type of $BM(\lambda)$ is the same as*

$$S_w^- A_\lambda S_{w'}^+.$$

Proof. Both maps depend only on the homotopy class of the path λ, so assume that λ is smoothly embedded in Y. Pick a Heegaard diagram $\mathcal{H}_w = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w} \cup \{w\})$ such that λ is embedded in Σ, and such that w' is in the same domain as $\{w\}$ and such that λ is entirely contained in this domain. Let $\mathcal{H}_{w'}$ denote $\mathcal{H}_w = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w} \cup \{w'\})$. There is an obvious identification of the chain complexes $CF^\circ(\mathcal{H}_w)$ and $CF^\circ(\mathcal{H}_{w'})$. Under this identification, the map $BM(\lambda)$ is just the identity. On the other hand, if we let $\mathcal{H}_1$ and $\mathcal{H}_2$ be the diagrams in Figure 21, and assume that λ is just the horizontal line between w and w' in those pictures, we have that

$$S_w^- A_\lambda S_{w'}^+ = S_w^- \Phi_{\mathcal{H}_1 \to \mathcal{H}_2} A_\lambda S_{w'}^+ = \begin{pmatrix} 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ * & 1 \end{pmatrix} \begin{pmatrix} 0 & U_w \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = 1,$$

using Theorem 14.1 and Lemma 14.6. It follows that $BM(\lambda) = S_w^- A_\lambda S_{w'}^+$. □

As a corollary, we can now prove that the path and graph cobordisms maps agree for graphs which are paths:

Corollary 14.8. *If $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a graph cobordism with Γ a collection of paths from the incoming to outgoing boundaries, then the path cobordism maps agree with the graph cobordism maps.*

Proof. We first pick a collar neighborhood N of Y_1 and an identification of N as $Y_1 \times I$. By choosing N appropriately we can assume that $N \cap \Gamma$ is a collection of paths from points in $Y_1 \times \{0\}$ to points in $Y_1 \times \{1\}$. Using Lemma 12.4 and the composition law for path cobordisms (Lemma 10.3), it is sufficient to show that the path and graph cobordisms agree for the cobordism $(N, \Gamma \cap N)$. Letting Γ_{Y_1} denote the projection of $\Gamma \cap N$ into Y_1, we note that the path cobordism map of $(N, \Gamma \cap N)$ is equal to the basepoint moving map $BM(\Gamma_{Y_1})$. Using Lemma 12.3, we know the graph cobordism map for $(N, \Gamma \cap N)$ is equal to the graph action map for Γ_{Y_1}. By Lemma 14.7, the graph action map in this case is equal to the basepoint moving map, and hence equal to the path cobordism map, completing the proof. □

As an additional consequence, we can now compute the $\mathbb{Z}_2[U_{\mathfrak{P}}]$-chain homotopy type of the π_1-action on $CF^\circ(Y, \mathbf{w}, \mathfrak{s})$. We first introduce some notation. Suppose that $w \in \mathbf{w}$ and pick a Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ for $(Y, \mathbf{w})$. Write the uncolored differential as

$$\partial = \sum_{n=0}^{\infty} \partial^n U_w^n,$$

where ∂^n has no U_w terms. We let Φ_w denote

$$\Phi_w = \sum_{n=1}^{\infty} n \partial^n U_w^{n-1},$$

the formal derivative of ∂ with respect to U_w. The map Φ_w was introduced by Sarkar in [Sar11b] when studying certain automorphisms of link Floer homology. The map Φ_w is a $\mathbb{Z}_2[U_{\mathbf{w}}]$ homomorphism and hence induces a $\mathbb{Z}_2[U_{\mathfrak{P}}]$-homomorphism on the colored complexes. Note however that the formal derivative $\frac{d}{dU_w}$ must be taken before coloring since $\frac{d}{dU_w}$ is not $\mathbb{Z}_2[U_{\mathbf{w}}]$-module homomorphism and hence doesn't induce a map on the tensor product of the uncolored CF° with the coloring module.

Remark 14.9. *If $w \in \mathbf{w}$, it is easy to see that Φ_w induces a map which is well defined on $CF^\circ(Y, \mathbf{w}, \mathfrak{s})$, the equivariant chain homotopy type invariant. This is observed because the transition maps commute with ∂, so upon applying d/dU_w to*

$$\partial \Phi_{\mathcal{H}_1 \to \mathcal{H}_2} + \Phi_{\mathcal{H}_1 \to \mathcal{H}_2} \partial = 0$$

and applying the Leibniz rule, we get that Φ_w commutes with the transition maps up to chain homotopy.

In the case that we are working with more than one basepoint in the component containing w, we also observe that on $CF^\circ(Y, \mathbf{w} \cup \{w\}, \mathfrak{s})$, the maps $S_w^+ S_w^-$ and Φ_w are chain homotopic. This follows since we simply pick another basepoint w' in the same component as w, and then pick a diagram so that

$$\partial = \begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix}$$

which can also be used to compute $S_w^\pm$. Differentiating with respect to U_w yields

$$\Phi_w = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix},$$

which is $S_w^+ S_w^-$.

At this point, we could prove the formula for the π_1 action by picking diagrams and multiplying a long string of matrices. While this is not hard, it is neither particularly elegant nor revealing. Instead we will try to prove things in a more invariant fashion. Surprisingly, we only need the computation from Theorem 14.1 so far as to see that the basepoint moving map is chain homotopic to $S_w^- A_\lambda S_{w'}^+$, which doesn't even require us to know the lower left component of $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ (which is Φ_w).

We start with a short lemma. The reader is encouraged to compare the first claim in the following lemma to the key computation we needed in [Zem15, Lem 7.3; Fig. 20]:

Lemma 14.10. *If λ is a path from w to w', we have that*

$$A_\lambda S_w^+ S_w^- + S_w^+ S_w^- A_\lambda \simeq 1.$$

Similarly if w and w' are basepoints which are in the same component of Y, we have

$$\Phi_w S_{w'}^\pm + S_{w'}^\pm \Phi_w \simeq 0.$$

Proof. For the first equality, take the expression

$$A_\lambda \partial + \partial A_\lambda = U_w + U_{w'},$$

and differentiate with respect to U_w. Using the Leibniz rule, we get that

$$A'_\lambda \partial + A_\lambda \partial' + A_\lambda \partial \frac{d}{dU_w} + \partial' A_\lambda + \partial A'_\lambda + \partial A_\lambda \frac{d}{dU_w} = 1 + (U_w + U_{w'}) \frac{d}{dU_w}.$$

Using the fact that $A_\lambda \partial + \partial A_\lambda = U_w + U_{w'}$ and the fact that $\partial' = \Phi_w$, we get exactly that

$$A_\lambda \Phi_w + \Phi_w A_\lambda \simeq 1.$$

As in Remark 14.9, we know that $\Phi_w \simeq S_w^+ S_w^-$, and hence we immediately get the first claim.

The second claim follows similarly. We just note that

$$\partial S_{w'}^\pm + S_{w'}^\pm \partial = 0,$$

and upon differentiating with respect to U_w, we get the desired claim. □

We can now prove Theorem F:

Theorem F. *If γ is a closed loop in Y starting and ending at w, the map $BM(\gamma)$ has $\mathbb{Z}_2[U_{\mathfrak{P}}]$ chain homotopy type*

$$BM(\gamma) = 1 + \Phi_w[\gamma].$$

Proof. Break γ into the concatenation of two paths, writing $\gamma = \lambda_2 * \lambda_1$. Say that λ_1 is a path from w to w' and λ_2 is a path from w' to w. Using Lemma 14.7, we then have that

$$BM(\gamma) = BM(\lambda_2) \circ BM(\lambda_1) = S_{w'}^- A_{\lambda_2} S_w^+ S_w^- A_{\lambda_1} S_{w'}^+.$$

Recall from Remark 14.9 that $S_w^+ S_w^- = \Phi_w$. Hence we have that

$$\begin{aligned}
BM(\gamma) &= S_{w'}^- A_{\lambda_2} \Phi_w A_{\lambda_1} S_{w'}^+ \\
&= S_{w'}^- A_{\lambda_2} (A_{\lambda_1} \Phi_w + 1) S_{w'}^+ && \text{(Lemma 14.10)} \\
&= S_{w'}^- A_{\lambda_2} A_{\lambda_1} \Phi_w S_{w'}^+ + S_{w'}^- A_{\lambda_2} S_{w'}^+ \\
&= S_{w'}^- A_{\lambda_2} A_{\lambda_1} S_{w'}^+ \Phi_w + S_{w'}^- A_{\lambda_2} S_{w'}^+ && \text{(Lemma 14.9)} \\
&= A_\gamma \Phi_w + S_{w'}^- A_{\lambda_2} S_{w'}^+ && \text{(Lemma 7.8)} \\
&= A_\gamma \Phi_w + 1 && \text{(Lemma 7.7).}
\end{aligned}$$

We finally note that $A_\gamma \Phi_w + \Phi_w A_\gamma \simeq 0$, which can be seen by differentiating $A_\gamma \partial + A_\gamma \partial = 0$. □

We now remark in applying $\frac{d}{dU_w}$ to an expression of the form $\partial(x)$, we must use the Leibniz rule, which yields

$$\frac{d}{dU_w}(\partial(x)) = \left(\frac{d}{dU_w}\partial\right)(x) + \partial\left(\frac{d}{dU_w}x\right)$$

or rather

$$\Phi_w = \frac{d}{dU_w} \circ \partial + \partial \circ \frac{d}{dU_w}.$$

The chain homotopy persists on any coloring where w is not given the same coloring as any other basepoint. We summarize with the following corollary:

Corollary 14.11. *The $\pi_1(Y, w)$ action vanishes on $HF^\circ(Y, \mathbf{w}, \mathfrak{s}, \sigma)$ for $\circ = +, -, \infty$ as long as w has a coloring distinct from all the other basepoints (in particular if $\mathbf{w} = \{w\}$). The chain homotopy is not $U_{\mathfrak{P}}$-equivariant, and hence the $\pi_1(Y, w)$ action does not necessarily vanish on $CF^\circ(Y, \mathbf{w}, \mathfrak{s}, \sigma)$, the equivariant chain homotopy type invariant.*

We can finally prove Theorem G:

Theorem G. *If $(W,\boldsymbol{\gamma}):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ is a path cobordism and each path is given a distinct coloring (for example if there is only one path) then the map $F^{\circ}_{W,\boldsymbol{\gamma},\mathfrak{s}}$ is independent of the curves $\boldsymbol{\gamma}$ on homology for $\circ=+,-,\infty$.*

Proof. This follows from the composition law. Suppose $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$ are two collections of paths. Decompose W into $W_2\circ W_1$, where W_1 is obtained by adding 1-handles and W_2 is obtained by adding 2- and 3-handles. Write $\partial W_1=-Y_1\sqcup Y'$. Observe that $\pi_1(Y')\to\pi_1(W)$ is surjective, so we can homotope $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$ so that they only differ in a collar neighborhood of Y'. Since W is a 4-manifold, such a homotopy can be taken to be an isotopy. Furthermore we can also assume that $\boldsymbol{\gamma}\cap W_i$ and $\boldsymbol{\gamma}'\cap W_i$ are paths for $i=1,2$. Using the composition law and the fact the π_1-action is trivial on homology for manifolds where all basepoints have distinct colors. □

The map Φ_w is not necessarily $U_{\mathbf{w}}$-equivariantly chain homotopic to zero, though we do have the following (cf. [Sar11b, Lem. 4.6]).

Proposition 14.12. *The map $(\Phi_w)^2$ is $U_{\mathbf{w}}$-equivariantly chain homotopic to zero.*

Proof. For this we need to work with $\mathbb{Z}$-coefficients, though we do not need naturality with $\mathbb{Z}$-coefficients, just the existence of a lift of the differential ∂ to $\mathbb{Z}$ coefficients for a given diagram. We start with

$$\partial\partial=0.$$

Differentiating with respect to U_w, we get

$$\Phi_w\partial+\partial\Phi_w+\partial\partial\left(\frac{d}{dU_w}\right)=\Phi_w\partial+\partial\Phi_w=0.$$

Differentiating again with respect to U_w, we get

$$\Phi_w'\partial+\Phi_w\Phi_w+\Phi_w\partial\left(\frac{d}{dU_w}\right)+\Phi_w\Phi_w+\partial\Phi_w'+\partial\Phi_w\left(\frac{d}{dU_w}\right)=0,$$

where Φ_w' is the formal U_w derivative of Φ_w. Using the fact that $\partial\Phi_w+\Phi_w\partial=0$ and doing some rearranging, we get that

$$\Phi_w\Phi_w=\frac{1}{2}\left(\Phi_w'\partial+\partial\Phi_w'\right)$$

and hence

$$\Phi_w\Phi_w=-\left(\sum_{n=2}^{\infty}\frac{n(n-1)}{2}\partial^n U_w^{n-2}\right)\partial-\partial\left(\sum_{n=2}^{\infty}\frac{n(n-1)}{2}\partial^n U_w^{n-2}\right).$$

Projecting to $\mathbb{Z}_2$ yields the desired chain homotopy for the chain complexes over $\mathbb{Z}_2$. □

References

[AE11] A. Alishahi and E. Eftekhary, *A refinement of sutured Floer homology* (2011), arXiv:1112.3540 [math.GT].

[BEH+03] F Bourgeois, Y Eliashberg, H Hofer, K Wysocki, and E Zehnder, *Compactness results in Symplectic Field Theory*, Geometry and Topology **7** (2003), 799-888.

[CGH10] V. Colin, P. Ghiggini, and K. Honda, *Embedded contact homology and open book decompositions*, arXiv:1008.2734[math.SG] (2010).

[CGH12a] ———, *The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions I*, arXiv:1208.1074[math.GT] (2012).

[CGH12b] ———, *The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions II*, arXiv:1208.1077[math.GT] (2012).

[CGH12c] ———, *The equivalence of Heegaard Floer homology and embedded contact homology III: from hat to plus*, arXiv:1208.1526[math.GT] (2012).

[GW08] J. E. Grigsby and S. M. Wehrli, *On the colored Jones polynomial, sutured Floer homology, and knot Floer homology*, arXiv:0807.1432[math.GT] (2008).

[Hir76] M. W. Hirsch, *Differential Topology*, Graduate Texts in Mathematics, Springer-Verlag, 1976.

[HT09] M. Hutchings and C. H. Taubes, *The Weinstein conjecture for stable Hamiltonian structures*, Geometry and Topology **13** (2009,) 901-941.

[Juh06a] A. Juhász, *Holomorphic discs and sutured manifolds*, Algeb. Geom. Topol. **6** (2006,) 1429-1457.

[Juh06b] ———, *Floer homology and surface decompositions* (2006), arXiv:0609779[math.GT].

[Juh15] ———, *Cobordisms of sutured manifolds and the functoriality of link Floer homology* (2015), arXiv:0910.4382v4[math.GT].
[Juh13] ———, *A survey of Heegaard Floer homology* (2013), arXiv:1310.3418v3[math.GT].
[Juh14] ———, *Defining and classifying TQFTs via surgery* (2014), arXiv:1408.0668[math.GT].
[JT12] A. Juhász and D. P. Thurston, *Naturality and mapping class groups in Heegaard Floer homology* (2012), arXiv:1210.4996 [math.GT].
[KM07] P. Kronheimer and T. Mrowka, *Monopoles and Three-Manifolds*, New Mathematical Monographs, vol. 10, Cambridge University Press, Cambridge, 2007.
[KLT10a] C. Kutluhan, Y. Lee, and C. Taubes, *HF=HM I: Heegaard Floer homology and Seiberg-Witten Floer homology* (2010), arXiv:1007.1979 [math.GT].
[KLT10b] ———, *HF=HM II: Reeb orbits and holomorphic curves for the ech/Heegaard-Floer correspondence* (2010), arXiv:1008.1595 [math.GT].
[KLT10c] ———, *HF=HM III: Holomorphic curves and the differential for the ech/Heegaard Floer correspondence* (2010), arXiv:1010.3456 [math.GT].
[KLT11] ———, *HF=HM IV: Seiberg-Witten Floer homology and ech correspondence* (2011), arXiv:1107-2297 [math.GT].
[KLT12] ———, *HF=HM V: Seiberg-Witten floer homology and handle additions* (2012), arXiv:1204.0115 [math.GT].
[Lip06] R. Lipshitz, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. **10** (2006), 955-1097.
[MW95] M. J. Micallef and B. White, *The structure of branch points in minimal surfaces in pseudoholomorphic curves*, Ann. of Math (2) **141** (1995), 35-85.
[MO10] C. Manolescu and P. Ozsváth, *Heegaard Floer homology and integer surgeries on links* (2010), arXiv:1011.1317v3 [math.GT].
[Ni10] Y. Ni, *Homological actions on sutured Floer homology* (2010), arXiv:1010.2808 [math.GT].
[OS04a] P. S. Ozsváth and Z. Szabó, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math (2) **159(3)** (2004), 1159-1245.
[OS04b] ———, *Holomorphic disks and topological invariants for closed three manifolds*, Ann. of Math. (2) **159** (2004), no. 3, 1027-1158.
[OS06] ———, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. in Math. **202,** (2006), Issue 2, 326-400.
[OS08] ———, *Holomorphic disks, link invariants, and the multi-variable Alexander polynomial*, Alg. & Geom. Top. **8** (2008), 615-692.
[Sar11a] S. Sarkar, *Maslov index formulas for Whitney n-gons*, J. of Symp. Geom. **9** (2011), Number 2, 251-270.
[Sar11b] ———, *Moving basepoints and the induces automorphisms of link Floer homology* (2011), arXiv:1109:2168 [math.GT].
[Tau07] C. H. Taubes, *The Seiberg-Witten equations and the Weinstein conjecture*, Geom. Topol. **11** (2007), 2117-2202.
[Zem15] I. Zemke, *A graph TQFT for hat Heegaard Floer homology* (2015), arXiv:1010.2808 [math.GT].

Department of Mathematics, UCLA, 520 Portola Plaza, Los Angeles, CA 90025, USA
E-mail address: ianzemke@math.ucla.edu